

AD732309

AFOSR - TR - 71 - 2792

Report No. 2188
Job No. 11432

INFORMATION PROCESSING MODELS AND
COMPUTER AIDS FOR HUMAN PERFORMANCE

FINAL REPORT, SECTION 4

Task 4: STUDIES OF HUMAN MEMORY
AND LANGUAGE PROCESSING

30 June 1971



ARPA ORDER NO. 890, Amendment No. 5

Sponsored by the Advanced Research Projects Agency,
Department of Defense, under Air Force Office of
Scientific Research Contract F44620-67-C-0033

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

Prepared for:

Air Force Office of Scientific Research
1400 Wilson Boulevard
Arlington, Virginia 22209

Reproduced by
NATIONAL TECHNICAL
INFORMATION SERVICE
Springfield, Va. 22151

DISCLAIMER NOTICE

THIS DOCUMENT IS THE BEST
QUALITY AVAILABLE.

COPY FURNISHED CONTAINED
A SIGNIFICANT NUMBER OF
PAGES WHICH DO NOT
REPRODUCE LEGIBLY.

Report No. 2188

Bolt Beranek and Newman inc.

INFORMATION PROCESSING MODELS AND
COMPUTER AIDS FOR HUMAN PERFORMANCE

FINAL REPORT, SECTION 4
Task 4: STUDIES OF HUMAN MEMORY
AND LANGUAGE PROCESSING
30 June 1971

by
Allan M. Collins
M. Ross Quillian

ARPA Order No. 890, Amendment No. 5
Sponsored by the Advanced Research Projects Agency,
Department of Defense, under Air Force Office of
Scientific Research Contract F44620-67-C-0033

*Pageinated by NTIS
in condition received*

Prepared for
Air Force Office of Scientific Research
1400 Wilson Boulevard
Arlington, Virginia 22209

Approved for public release;
distribution unlimited.

TABLE OF CONTENTS

	<u>Page</u>
SUMMARY.....	iii-v
1. PREFACE.....	1
2. ANNOTATED BIBLIOGRAPHY OF PAPERS.....	2
3. OVERVIEW.....	5
4. REPORTS.....	11

Report No. 2188

Bolt Beranek and Newman Inc.

FINAL TECHNICAL REPORT

ARPA Order No. 890, Amendment No. 5

Program Code No. 9D20

Contractor: Bolt Beranek and Newman Inc.

Effective Date of Contract: 1 November 1966

Contract Expiration Date: 30 June 1971

Amount of Contract: \$804,896.00

Contract No. F44620-67-C-0033

Principal Investigators: John A. Swets

Mario C. Grignetti

Wallace Feurzeig

M. Ross Quillian

Telephone No. 617-491-1850

Title: INFORMATION PROCESSING MODELS AND
COMPUTER AIDS FOR HUMAN PERFORMANCE

TASK 4: STUDIES OF HUMAN MEMORY AND LANGUAGE PROCESSING

1. Technical Problem

The aim of this project was to determine how people store and retrieve factual (non-numerical) information and how they utilize this stored information in comprehending English text. Three of the studies investigated how people retrieve factual information, one evaluated possible strategies for interpreting text, and a final paper summarized our conclusions about the requirements for building a computer-based, natural-language-processing system.

2. General Methodology

Laboratory experiments.

3. Technical Results

The results have indicated that people use both deductive inference and inference by analogy in answering questions. The initial search for relevant facts is apparently a parallel process, while the checking of possible answers is a serial process. Depending on the information turned up by the parallel search and the constraints of syntax and context, people apply a variety of different specific decision rules in order to decide how to answer a question or how to interpret a sentence.

4. Department of Defense Implications

Military operations in the future will utilize computer-based, question-answering systems that can store and retrieve factual

information and that can interact with users in English. Knowledge gained from these experiments is being used in a computer project aimed toward developing such systems.

5. Reports Annotated Within

Collins, A. M. and Quillian, M. R. Facilitating retrieval from semantic memory: The effect of repeating part of an inference. Attention and Performance III (ed. by A. F. Sanders), Acta Psychologica, 1970, 33, 304-314.

Collins, A. M. and Quillian, M. R. Does category size affect categorization time? Journal of Verbal Learning and Verbal Behavior, 1970, 9, 432-438.

Collins, A. M. and Quillian, M. R. Tripping down the Garden Path.

Collins, A. M. and Quillian, M. R. Categories and Subcategories in Semantic Memory.

Collins, A. M. and Quillian, M. R. How to make a language user. To be published in Organization and Memory (ed. by E. Tulving) New York: Academic Press, 1972.

UNCLASSIFIED

Security Classification

DOCUMENT CONTROL DATA - R & D

(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1. ORIGINATING ACTIVITY (Corporate author) Bolt Beranek and Newman Inc. 50 Moulton Street Cambridge, Massachusetts 02138		2a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED	
		2b. GROUP	
3. REPORT TITLE INFORMATION PROCESSING MODELS AND COMPUTER AIDES FOR HUMAN PERFORMANCE TASK 4: STUDIES OF HUMAN MEMORY AND LANGUAGE PROCESSING			
4. DESCRIPTIVE NOTES (Type of report and inclusive dates) Scientific Final			
5. AUTHOR(S) (First name, middle initial, last name) Allan M. Collins M. Ross Quillian			
6. REPORT DATE 30 June 1971	7a. TOTAL NO. OF PAGES 118	7b. NO. OF REFS 56	
8a. CONTRACT OR GRANT NO F44620-67-C-0033		8b. ORIGINATOR'S REPORT NUMBER(S)	
b. PROJECT NO 890			
c. 61101D			
d. 681313		9b. OTHER REPORT NO(S) (Any other numbers that may be assigned this report) AFOSR - TR - 71-2792	
10. DISTRIBUTION STATEMENT Approved for public release; distribution unlimited.			
11. SUPPLEMENTARY NOTES TECH, OTHER		12. SPONSORING MILITARY ACTIVITY AIR FORCE OFFICE OF SCIENTIFIC RESEARCH (NL) 1400 WILSON BLVD ARLINGTON, VIRGINIA 22209	

13. ABSTRACT

The aim of this project was to determine how people store and retrieve factual (non-numerical) information and how they utilize this stored information in comprehending English text. Three of the studies investigated how people retrieve factual information, one evaluated possible strategies for interpreting text, and a final paper summarized our conclusions about the requirements for building a computer-based, natural language-processing system. The results have indicated that people use both deductive inference and inference by analogy in answering questions. The initial search for relevant facts is apparently a parallel process, while the checking of possible answers is a serial process. Depending on the information turned up by the parallel search and the constraints of syntax and context, people apply a variety of different specific decision rules in order to decide how to answer a question or how to interpret a sentence. Military operations in the future will utilize computer-based, question-answering systems that can store and retrieve factual information and that can interact with users in English. Knowledge gained from these experiments is being used in a computer project aimed toward developing such systems.

DD FORM 1473

NOV 68

UNCLASSIFIED

Security Classification

1. PREFACE

At its inception in 1966, this contract was devoted solely to the one area of second-language learning. Later amendments have added three more tasks: Models of Man-Computer Interaction; Programming Languages as a Tool for Cognitive Research; and Studies of Human Memory and Language Processing. The present contract was scheduled for termination on 31 December 1970, but the final reporting date was changed to 30 June 1971, to allow completion of data analysis in the various tasks.

Due to the amount of information to be presented in the Final Report, we have bound it in four Sections, one for each task. In addition to a copy of this page, each Section contains an appropriate subset of the documentation data required for the report: a contract-information page, a summary sheet for the particular task at hand, and a DD form 1473 for document control.

2. ANNOTATED BIBLIOGRAPHY OF PAPERS PRODUCED FOR THIS PROJECT

Collins, A. M. and Quillian, M. R. Facilitating retrieval from semantic memory: The effect of repeating part of an inference. Attention and Performance III (ed. by A. F. Sanders), Acta Psychologica, 1970, 33, 304-314.

This experiment tested the hypothesis that people decide whether a sentence like "A canary can fly" is true or false by inference from the two facts that a canary is a bird and that birds can fly. This hypothesis has an implication for reaction time (RT) in deciding about pairs of such sentences presented in succession. Prior exposure to one sentence should reduce RT to a second sentence if the same fact is involved in confirming both. This prediction held for the eight different conditions in which it was tested.

Collins, A. M. and Quillian, M. R. Does category size affect categorization time? Journal of Verbal Learning and Verbal Behavior, 1970, 9, 432-438.

Two experiments were run to determine why it takes people longer to categorize object names (e.g., collie or tulip) into larger categories (e.g., animal) than into smaller categories (e.g., dog). It was found that this difference was due to the nesting of the smaller categories in the larger categories, and not to the difference in category

size. It was argued that categorization time for "No" responses (e.g., tulip) depends on how closely related in terms of semantic distance the given category (e.g., animal) is to the correct category (i.e., plant).

Collins, A. M. and Quillian, M. R. Tripping Down the Garden Path.

Two experiments were run to see how people revise a mistaken interpretation in part of a sentence. A reaction-time task was used where ss decided whether a string of words was a sentence or not. Among the sentences were some that were likely to be misinterpreted at first (i.e., garden-path sentences). The results showed that garden-path sentences take longer to interpret than normal sentences but that this effect is independent of the number of words in the sentences. Apparently, reprocessing in the garden-path sentences only involved those words that were misinterpreted initially.

Collins, A. M. and Quillian, M. R. Categories and Subcategories in Semantic Memory.

Subjects categorized names of animals and plants with respect to three different categories: "animal," "bird," and "mammal." There were four kinds of lists: one kind contained only animals that were mammals, a second kind both mammals and non-mammals, a third kind only birds, and a fourth kind both birds and non-birds. The results indicated that the category

mammal is not intermediate between elephant and animal, in the way that bird is intermediate between robin and animal, since it takes longer to decide that an elephant is a mammal than an animal, but less time to decide that a robin is a bird than an animal. The results also showed it takes longer to decide about a robin or an elephant when there are non-birds or non-mammals included in the list.

Collins A. M. and Quillian, M. R. How to make a language user. To be published in Organization and Memory (ed. by E. Tulving) New York: Academic Press, 1972.

This paper provides a top-level description of what we think is required to build a computer-based, natural-language-processing system that can comprehend text, store information, and retrieve answers to questions in the same way that people perform these operations. This paper summarizes in an integrated manner most of the knowledge we have accumulated during this project.

3. OVERVIEW

Computerized question-answering systems that converse in English will probably be used for storing and retrieving military information in the not-too-distant future. In this project, we have conducted experiments on how humans perform aspects of these tasks in order to aid the development of such computer systems. These findings are being utilized in a computer project that is developing a network for storage of factual information and routines for conversing with this network in English.

There are three general advantages for the development of computerized question-answering systems that derive from these psychological experiments: (a) knowing how people process natural-language information provides strategies for computer programs to do the same processing (programmers now try to analyze their own processing introspectively, which is quite unreliable); (b) accessing information by its "associative" semantic structure, as humans do, will make it unnecessary to anticipate with an indexing scheme how the information will be requested in the future; and (c) knowledge of human information processing will guide development toward systems that interact with man in the most efficient way.

We will briefly summarize here the important conclusions we have reached about human semantic memory that have implications for building a computer-based, natural-language-processing system. These points are discussed at length in the paper "How to make a language user." The first set will concern what is stored in human memory and the second set how that information is processed.

What is stored:

(1) Much of what people know (e.g., that Aristotle could talk) is never learned or stored *per se*. Instead it is inferred from what is stored; in the example, the inference follows from the fact that Aristotle was a man and men talk. If one considers all the properties known about people, and all the people one knows, then it becomes evident how sizeable is the economy of storage gained by not storing each property with each person directly. This kind of economy applies everywhere in human memory.

(2) Most information is not stored in quantified form. Thus, a person usually does not store whether all birds or most birds can fly or have wings. If such information is needed, people search memory for examples of birds that do not fly or do not have wings. All estimates of what proportion of things have a given property (what proportion of birds can sing) are based on a search for positive and negative examples, and an evaluation is based on the numbers of each type found.

(3) People store negative facts (e.g., "A penguin can't fly") only when the information contradicts something that might be inferred by mistake or something that is true for similar concepts. People do not store information like "Ships don't have wings" but must infer such contradictions when needed by methods described below in the set of memory processing.

(4) New concepts are set up in memory in cases such as "young dog" or "South American countries" whenever information is learned that cannot be derived from the descriptive label itself (e.g., young dogs are called puppies and are frisky; South American countries are Spanish speaking, except for Brazil). The phrases "brown dog" or "coastal city" are examples of phrases that are not concepts.

(5) Concepts often have more than one superordinate directly stored. Hence Bolivia is both a country and a South American country; an eagle is a bird and a bird-of-prey. Even though some superordinates are stored directly, others are only reached by going up the chain of superordinates. Hence, an eagle is also an animal, a living thing, and an object, but these must be inferred from the knowledge that birds are animals, animals are living things, and living things are objects. As suggested in (1), all the properties that hold for any superordinate of a concept also hold for the concept itself, unless the negation is stored directly with the concept as in (3).

(6) There are other special relations which, like superordinate, permit whole classes of inferences to be made. The major examples of these are: similarity, part, proximity, consequence, precedence, parent of. For example, to know that Katmandu is part of Nepal permits one to infer information about its location, its climate, its topography, and its maximum size in area and population; that is, assuming one knows such information about Nepal.

(7) These special relations form the bases for grouping or organizing concepts in memory. Often several of these relations apply to the same subject matter so that there are overlaying organizations. For example, in anatomy hands and feet are grouped on the bases of similarity and neck and shoulders on the bases of proximity. The reason why organization occurs in memory is so that inferences can be made; i.e., so that people can use their memories in a generative manner.

How information is processed:

(8) Both comprehension and retrieval involve a semantic search in memory for paths or connections between concepts. This search goes out in parallel from all the words in the sentence or question spreading out from each concept to all directly related concepts. When the search originating from one word encounters the search from another, a connection or path between the two concepts has been found through other concepts. When a connection is found, an interrupt occurs, and the connection is checked to see if it meets the constraints of syntax and context.

(9) In language processing, much information is processed tacitly in parallel as described above, but never explicitly unless there is something that causes an interrupt. For example, the sentence "The policeman held up his hands to stop the cars" does not produce explicit processing of the fact that people are pushing the brakes in the cars referred to. But, if told previously that an earthquake had started the

cars rolling down the hill, readers of the above sentence would wonder how the policeman could stop the cars. In other words, what is tacitly processed in one context is explicitly processed when there is some anomaly that causes an interrupt.

(10) Many aspects of language processing involve making decisions as to whether two concepts are equivalent within the constraints of syntax and context. The question of equivalence arises in dealing with nouns and pronouns that refer to earlier words in text, in dealing with metaphor, and in dealing with simple questions such as "Does a canary quack"? or "Is a stagecoach a vehicle? In short, it arises in every aspect of language processing.

(11) There are a number of different decision strategies or decision rules that people use to decide whether two concepts are equivalent. The decision rules depend both on the connections found and the constraints of syntax and context. These rules are cited in the paper "How to make a language user." The decision strategies for rejecting equivalence of two concepts depend on finding a connection that leads to contradiction of some kind.

(12) In storing information, the use of language causes properties that are common to different examples to be stored with higher-level concepts and the distinguishing properties of each example to be stored with lower-level concepts. For instance, if a vulture, cardinal, and canary are all referred as birds, then the kind of

semantic search described will find connections through properties they have in common, and these will be stored with the concept bird. If instead they are referred to as vultures, cardinals, and canaries, then the semantic search will find the distinguishing properties and these will be stored with concepts like vulture, canary, or cardinal.

While these twelve points only briefly touch on much of what we have learned that is relevant to building natural-language-processing systems, they do summarize the kinds of ideas we plan to implement in computer systems in the near future.

4. REPORTS

The last two papers annotated above are included in this report immediately after this page. The first three papers were included in earlier reports.

BLANK PAGE

CATEGORIES AND SUBCATEGORIES
IN SEMANTIC MEMORY*

Allan M. Collins
M. Ross Quillian

Bolt Beranek and Newman Inc.
50 Moulton St.
Cambridge, Mass. 02138

*This research was supported by the Advanced Research Projects
Agency of the Department of Defense and was monitored by the
Air Force Office of Scientific Research under Contract No.
F44620-67-C-0033.

ABSTRACT

Subjects categorized names of animals and plants with respect to three different categories: "animal," "bird," and "mammal." There were four kinds of lists: one kind contained only animals that were mammals, a second kind both mammals and non-mammals, a third kind only birds, and a fourth kind both birds and non-birds. In each list there was an animal name that was semantically similar to the first animal name in the list and one that was unrelated to the first animal name.

The results indicated that the category mammal is *not* intermediate between elephant and animal in the way that bird is intermediate between robin and animal, since it takes longer to decide that an elephant is a mammal than an animal, but less time to decide that a robin is a bird than an animal. The results also showed that it takes longer to decide about a robin or an elephant when there are non-birds or non-mammals included in the list. Lastly, it was found that semantic similarity had quite different effects on decision time depending on whether the correct response was "Yes" or "No."

INTRODUCTION

There are categories or groupings of concepts that are learned very early in the semantic development of children, such as dogs, birds, animals, cars, boats, and even collies and eagles. On the other hand, there are categories learned later, such as birds-of-prey, canines, mammals, farm animals, vehicles and warships, which must somehow be added to the semantic structure already built. It is possible to envision at least two rather tidy schemes for representing the inclusion relationships between categories, but we have argued (Collins and Quillian, in press) that, because of the differences in the way categories are learned, the resulting structure is not at all tidy. Furthermore, it can be misleading *not* to pay attention to the irregularities of structure in designing experiments on semantic memory.

One scheme for structuring memory, as proposed by Kintsch (in press) is based on pointers between concepts (in this respect it is like our model). In his scheme a higher-order concept such as animal can be partitioned in different ways: for example, (1) pet V farm animal V wild animal, (2) mammal V bird V insect V reptile, etc., (3) human V non-human. Then categories like elephant might map into each of these different partitions; that is, an elephant would have pointers to wild animal, mammal, and non-human. Given such a view, the most likely assumption about processing implies the following: to decide an elephant is an animal, it is necessary to reach animal via the path through wild animal, mammal, or non-human. Hence, it should take longer to decide an elephant is an animal than a mammal, unless one makes the rather implausible assumption that the paths to animal via wild animal or non-human are shorter

than the path via mammal by a factor of two. The factor of two occurs because to reach mammal it is only necessary to travel half way along the latter path. Though Kintsch may not hold to it rigidly, this kind of view assumes that categories learned later are inserted into the structure in a way that preserves a partial ordering of concepts from lower-order to higher-order concepts.

A comparable, but less stringent, assumption about structure is made by representing inclusion of concepts in terms of Venn diagrams, as Meyer (1970) does. As illustrated in Figure 1, even a simple two-dimensional Venn diagram provides a fairly flexible way to represent inclusion relationships between concepts. As in Figure 1, the diagram can be drawn so that sponge is closer to plant than elephant is, and man can be shown as partly animal and partly not. Venn diagrams, which can be extended to n-dimensions, also correspond roughly to physiological theories of memory which make reference to fields (Lashley, 1949), cell assemblies (Hebb, 1949), or foci (John, 1966). But even though Venn diagrams can sometimes be helpful in thinking about similarity and superordinate relations between concepts,

[Insert Figure 1 about here]

they can also be very misleading if we are correct. This is because the nature of Venn diagrams forces a concept like mammal to be intermediate between elephant and animal in the same way that bird is intermediate between robin and animal. The assumptions involved in using Venn diagrams are not as stringent as in the pointer model, because there is no implication that deciding an elephant is a mammal will take less time than




FIG.1 A VENN DIAGRAM REPRESENTATION OF INCLUSION AND DISTANCE RELATIONS BETWEEN CONCEPTS

BLANK PAGE

deciding it is an animal. In fact, one plausible processing hypothesis, that "yes" reaction time (RT) is inversely related to the distance of the concept from the edge of the category would predict the opposite result. That is to say, it might take longer to decide an elephant is a mammal than an animal because elephant is closer to the edge of the mammal concept than the animal concept. Even though use of Venn diagrams is non-committal about such processing considerations, still if processing leads to different time order relations as between elephant, mammal, and animal on the one hand and robin, bird, and animal on the other, then Venn diagrams are seriously misrepresenting the underlying structure.

In contrast to these two views, we (Collins and Quillian, in press) have argued that children first learn that a concept like elephant is a kind of animal, that a robin is a kind of bird, and that a bird is a kind of animal. When the concept mammal is learned later, there is no change made in the earlier structure. Instead, any facts about mammals, such as the fact that a dog is a mammal, are stored in addition. Hence, there will be pointers from only a few animals (e.g., whale, bat, kangaroo, dog, maybe elephant) to the concept mammal, and even for these cases, the concept mammal is likely to be a less accessible category than the concept animal. The implication of our view for this discussion is that it should take longer to decide a robin is an animal than a bird, because it is necessary to go through bird to get to animal. In contrast, it should take less time to decide an elephant is an animal than a mammal, because mammal is at best a secondary superordinate category for elephant. If this prediction is correct, then there are clear dangers in using partial orderings or Venn diagrams to represent inclusion relations between concepts.

To test these ideas, we used a categorization task where reaction time (RT) was measured for Ss to decide whether an instance such as elephant or robin presented on a display belonged to a prespecified category. The category remained constant for a list of fourteen trials, and there were eight such lists seen by each S, two with the category "bird," two with the category "mammal," and four with the category "animal." For each instance like elephant, one group of Ss decided whether it was a "mammal" and another group whether it was an "animal." Similarly, for robin, one group decided whether it was a "bird" and another group whether it was an "animal."

With this method, there were also two questions about processing we wanted to investigate. The first question arose from an earlier study of categorization (Collins and Quillian, 1970b) where we found people categorize names of dogs (e.g., collie) faster than names of birds (e.g., robin) or animals (e.g., elephant). A clue to why dogs were categorized faster was suggested by one S who indicated that she was surprised when she encountered lizard in the "animal" list after animals like beaver, elephant and goat. We refer to lizard as a *wide* instance for the category "animal" because it is outside the range of instances of animals (roughly mammals) that Ss expect when the category is "animal." The inclusion of wide instances in a list probably would slow Ss down at least for the wide instances themselves and perhaps also for narrow instances, in this case mammals. In the earlier study, we hypothesized that dogs were categorized faster, because of the lack of wide instances for the category dog. In this study, then, we decided to construct two kinds of animal lists, one kind with only mammals (*narrow* lists) and another kind with non-mammals as well, such

as crocodile, octopus, or frog (*wide lists*). Equivalent narrow and wide lists were constructed for both the "mammal" category and the "animal" category. It should be noted that a wide instance was a "No" response for the category "mammal" and a "Yes" response for the category "animal." The same distinction was carried over to the lists of birds by adding animals that were not birds (mostly mammals) to make wide lists.

Because we wanted to see whether or not exposure to a narrow list slowed down categorization of a wide instance, we included a wide instance near the end of each narrow list for comparison with wide instances at both the beginning and end of wide lists. Such a slowdown could occur if the S, after seeing only mammals when the category is "animal," restricts his effective category to some self-chosen category similar to "mammal." It would be somewhat paradoxical if such a self-chosen category similar to "mammal" produced *faster* RTs for narrow instances on a narrow "animal" list as opposed to wide "animal" lists, whereas the category "mammal" produced *slower* RTs for narrow instances on a "mammal" list as opposed to an equivalent "animal" list.

The other question about processing we wanted to investigate was whether Ss will utilize a decision made earlier during the list in deciding about a semantically similar instance. For example, suppose crocodile is one of the first instances in either an "animal" or a "mammal" list. Then several trials later alligator occurs. Our prediction was that the S should be faster in deciding about alligator than he would be for an equivalent unrelated instance such as octopus or frog. We expected such a "priming" effect would occur both for cases where the instance

(alligator) was in the category ("animal") and where it was not ("mammal"). A priming effect for alligator could occur either because the S can follow the semantic path in memory faster the second time (Collins and Quillian, 1970a) or because the prior response, "Yes" or "No," would be stored directly with crocodile, and the S would merely need to reach crocodile from alligator to find the correct response.

METHOD

The 16 Ss were employees of BBN. All were naive as to the nature of the experiment. The words were displayed one at a time on a cathode-ray tube (CRT) attached to a computer. The S sat about 3 ft away from the screen, and the words varied from about 4° to 8° visual angle on the screen. First, a warning dot came on the screen for 0.5 sec, followed by the category name in quotation marks for 2 sec. Then, there was a 0.5 sec pause followed by the warning dot for 0.5 sec and the word to be categorized for 2 sec. The word to be categorized, which we call an instance, was *not* in quotation marks. The same timing cycle repeated through all the trials. The S responded by pressing the right-hand microswitch if the word was in the category, and the left-hand microswitch if it was not. The S's response was recorded if it occurred anytime during the two seconds the word was on the screen.

There were three categories used: "animal," "mammal" and "bird." The category remained the same for a list of 14 trials in a row. Within each list, about half the words belonged to the category and half did not. Each S saw eight such lists: two with the category "mammal," two with the category "bird," and four with the category "animal." There were two different kinds of lists, which we call *wide* lists and *narrow* lists. The Ss knew nothing about this distinction. For the category "mammal," the wide list included both mammals, such as beaver, camel, and sheep (seven of these narrow instances), and animals that were not mammals, such as spider, alligator, and lobster (three of these wide instances). For the wide instances the correct response was "no." Both kinds of instances included only those animals that the Thorndike-Barnhart Beginning Dictionary (1968) defines as

animals, and not instances such as haddock, which is defined as a fish. The narrow list included only narrow instances (six of these), except that one of the last two words was a wide instance, i.e., an animal that was not a mammal. This was the same word as the first wide instance (one of the first two words) in the equivalent wide list. The "no" responses on both lists that were not wide instances were all plants and vegetables.

For both the wide and narrow "mammal" lists, there was an equivalent "animal" list. As with "mammal" lists, the narrow "animal" list included only animals that were mammals, except for the wide instance near the end of the list. The wide "animal" list included both animals that were mammals (three of these) and animals that were not (four of these). As before, the wide instance at the end of the narrow list was the same as the one at the beginning of the wide list. When the category was "animal," the wide instances were "yes" responses. In order to keep the number of "yes" and "no" responses equal in both wide lists, there were four mammals in the wide "mammal" list that were not in the equivalent wide "animal" list, and the data from these four dummy instances were ignored. Altogether, there were three mammals (narrow instances) that occurred in both wide and narrow "mammal" lists and in the equivalent wide and narrow "animal" lists.

When the category was "bird," there was also a distinction between wide and narrow lists. A narrow list included only animals that were birds, except for a wide instance at the end of the list. A wide list included both birds and other animals, mostly mammals that were wide instances for the "bird" lists. Just as before, there were wide and narrow "animal" lists equivalent to the wide and narrow "bird" lists, and there were three birds that occurred in the four equivalent "bird" and "animal" lists. Because

there were "animal" lists equivalent to both "mammal" and "bird" lists, we will refer to an "animal" list as a mammal "animal" list or a bird "animal" list when it is necessary to distinguish them.

In order to counterbalance words in the lists and have each S contribute RTs to all eight conditions (four kinds of lists for both birds and mammals), we divided the Ss into four different groups which saw the eight conditions in different orders. This involved constructing four different sets of equivalent lists with mammals in them and four with birds. The narrow and wide instances seen by each of the four groups, and the order in which they appeared are shown in the Appendix. The plants and vegetables that were used for "no" instances and the dummy "yes" instances added in the wide "mammal" and "bird" lists are omitted in the Appendix, but were inserted pseudorandomly in the actual lists presented. The important aspect of these different groups is that for almost every comparison made with the data, the same words were used, but they appeared in different conditions for different groups of Ss. For example, crocodile occurred at the beginning of a wide "animal" list for one group, at the end of a wide "animal" list for a second group, and at the end of a narrow "animal" list for a third group.

There was one other variable in the lists. The word at the beginning of each list, as shown in the Appendix, was semantically similar to one of the words in the middle of the list and semantically unrelated to another of the words (the control instance). For example, if squid was the first wide instance in a wide "mammal" or "animal" list, then a similar instance in the middle of the list might be octopus and the unrelated instance might be salamander. For a narrow "mammal" list, the first

instance might be beaver, the similar instance raccoon, and the unrelated instance leopard. For a narrow "bird" list, the first instance might be parakeet, the similar instance canary, and the unrelated instance goose. Here again, the words and orders were counterbalanced across groups. Thus, if octopus was the similar instance and salamander the unrelated instance for one group, for another group the first word was lizard so that salamander was the similar instance and octopus was the unrelated instance for this group.

RESULTS AND DISCUSSIONS

In analyzing the results, the means across Ss were computed for correct responses only. We used difference *t* tests to analyze the mean RTs for paired conditions averaged across all 16 Ss. In the experiment, approximately six percent of all the responses were errors or omissions.

The average RTs for the counterbalanced narrow instances are shown in Figure 2. The abscissa shows the different relative positions of the three narrow instances on the list. The actual positions on the lists varied from 3 to 4 for N1, 5 to 10 for N2, and 9-12 for N3.

Considering only the category distinctions, "bird" vs "animal" on the left and "mammal" vs "animal" on the right, it is evident that a bird name can be categorized as a "bird" faster than as an "animal," $t(5)=3.32, p<.05$, whereas a mammal name can be categorized as an "animal" faster than as a "mammal," $t(5)=3.44, p<.05$. The same pattern also holds later in Figure 4 for the data based on bird names and mammal names. Hence, the prediction that the category "mammal" and the category "bird" are related to the category "animal" in different ways was confirmed.

Insert Figure 2 about here

Considering the distinction between wide and narrow lists in Figure 2, there is a significant tendency in positions N2 and N3, for RT to narrow instances to be faster in a narrow list than in a wide list, $t(7)=7.9, p<.01$. As would be expected, the

difference between wide and narrow lists does not show up in the N1 position, because the Ss need to see several instances before they build up an expectation of what kind of instances will appear. The difference between narrow and wide lists occurs because Ss become faster at categorizing narrow instances as they go through a narrow list, but do not become faster on narrow instances as they go through a wide list.

There is one apparent anomaly in the first position (N1) between the wide and narrow "animal" lists on the left. This large difference occurred because Ss had already seen a bird name (e.g., robin) in the narrow "animal" list and had not in the wide "animal" list; hence in the former case on our theory, Ss had already made the inference once before that a bird is an animal, whereas Ss had not in the latter case. Thus, the difference is due to a facilitation effect from a previous inference, an effect we have found elsewhere, (Collins and Quillian, 1970a). The fact that the same difference did not occur for the two "animal" lists on the right is further evidence that there is no similar kind of inference involved in deciding an elephant, for example, is an animal.

The average RTs for the "No" instances that were plants and vegetables are shown in Table 1.

Insert Table 1 about here

The differences between conditions in Figure 2 are not apparent in these data. Hence, the "No" RTs for plants and vegetables appear to be largely independent of any manipulations of the kinds of animals shown in the lists.




FIG. 2 MEAN REACTION TIMES FOR NARROW INSTANCES AT DIFFERENT
RELATIVE POSITIONS IN NARROW AND WIDE LISTS

TABLE 1

Mean Reaction Times for "No" Instances (Plants and Vegetables)
In Different Conditions

	Category			
	"Bird"	"Animal"	"Mammal"	"Animal"
Narrow list	832	784	797	808
Wide list	851	787	845	811

In Figure 3, the average RTs for wide instances at the end of narrow lists and at different relative positions in wide lists are shown. The actual positions in the lists varied from 1 to 2 for W1, 5 to 8 for W2 and 7 to 10 for W3, and 13 to 14 for W4. The correct response for a wide instance was "No" when the category was "bird" or "mammal" and "Yes" when the category was "animal." The pattern of results is quite different (though not significantly different because of the paucity of data points) in the two cases. In all cases Ss can categorize wide instances faster as they go through a wide list. But RT for a wide instance at the end of a narrow list seems to depend on whether

Insert Figure 3 about here

the correct response is "Yes" or "No." For the "Yes" response, a wide instance at the end of a narrow list is categorized about as fast as one at the end of a wide list, and much faster than a wide instance at the beginning of a wide list. However, for the "No" responses, a wide instance at the end of a narrow list is categorized about as fast as a wide instance at the beginning of a wide list.

We certainly did not expect to find a difference between "Yes" and "No" responses, and in fact, we suspected that after seeing a whole list of narrow instances, the Ss might actually be slower in categorizing a wide instance than they are at the beginning of a list. Thus, at least two questions are raised by this result: (1) Why is there a difference between "Yes" and "No" responses and (2) For "Yes" responses, why are Ss faster in categorizing a wide instance at the end of a list of narrow instances than they would have been at the beginning

of the list? With respect to this latter question, there was no tendency to categorize narrow instances faster toward the end of a wide list (as is evident from Figure 2), even though a wide list contained several narrow instances. So rephrasing the second question, why was there improvement in this condition when Ss had no previous exposure to instances of the same type in the list, whereas there was no improvement in a condition where Ss did have previous exposure to instances of the same type in the list?

To provide a plausible answer to these two questions, we have to fall back on aspects of the task that confronted the Ss. Thus, our explanation is *ad hoc* and task dependent, and we offer it only to show that these results are not incompatible with our general theoretical framework. We think that the difference between the "Yes" and the "No" responses to wide instances at the end of the narrow lists has to do with the Ss forming a subjective category for the "No" responses that were plants and vegetables. By the end of a narrow list, and probably in the experiment as a whole, the Ss would learn to respond "No" whenever they see a a plant or vegetable name. In addition to a subjective "No" category for plants and vegetables, the Ss would also form a subjective "Yes" category in a narrow list roughly equivalent to either mammals or birds depending on the kind of list. Hence, in narrow lists, there would likely be two subjective categories apart from the given category.

Suppose the given category is "animal" and the S has seen a narrow list with only mammals, and plants and vegetables. When spider appears at the end of such a list, it does not fit within the subjective category, but a spider is an "animal" and so it can be categorized rather quickly into the category "animal."




FIG. 3 MEAN REACTION TIMES FOR WIDE INSTANCES AT DIFFERENT RELATIVE POSITIONS IN WIDE LISTS AND AT THE END OF NARROW LISTS

BLANK PAGE

The category "animal" will be more available at the end of a narrow list than at the beginning of a wide list, because the S has seen the category name 13 or 14 times in a row at the end of the list, whereas he has seen it only once or twice at the beginning of a list. Suppose on the other hand that the given category is "mammal," and again the S has seen a narrow list with only mammals, and plants and vegetables. When spider appears at the end of this list, it fits neither the two subjective categories, nor the given category "mammal." Because the subjective "No" category is the only "No" category, the S has available, he may be slowed down in double checking any inference with respect to the given category that leads to a "No" response. If the S is double checking inferences in this case, then the question again arises as to why he is not slower at the end of a narrow list than at the beginning of a wide list? Our answer is the same as for the "Yes" responses. At the end of a list he has worked with the category name "mammal" 13 or 14 times, and he can make any decision with respect to that category name faster at the end of a list than at the beginning, so that any double checking is a second factor that offsets the decrease in RT from working with the same category 13 or 14 times. The same logic applies among the bird lists used.

The final question then is why there is no decrease in RT for narrow instances in a wide list, if he becomes faster in working with the same category 13 or 14 times? First, notice in Figure 2 that Ss did become faster on bird names (narrow instances) in the wide list when the category is "animal"? The probable reason is that Ss adopted two subjective categories in the wide list, bird and animal, rather than making the inference each time that a bird is an animal. As a general strategy then,

Ss may have been adopting multiple categories to deal with wide lists. This can account for the very sharp decreases in RT for wide instances in a wide list. But what happens to narrow instances if the S adopts multiple categories?

If other categories are added, the increased number of categories may act to slow down RT to any one category, in particular the given category which applies to the narrow instances. In some sense, an increase in RT as the number of categories increased might parallel the increase, Sternberg (1966) or Nickerson (1966) find in memory search tasks with an increase in the number of targets. Such an increase in RT has recently been found by Juola and Atkinson (in press) who varied the number of target categories from one to four. The point here is that the increase in number of subjective categories for wide lists may act to offset the practice effect that can be seen in Figure 2 for narrow instances in a narrow list.

The other processing question we investigated in this experiment was the effect of semantically similar instances on the RT of instances seen later in the list. The average RTs for instances unrelated to any earlier instance and for instances similar to an earlier instance are shown in Figure 4. Our prediction was that the similar instances would be facilitated by the earlier instance. This appears to be true for the "No"

Insert Figure 4 about here

responses, though the difference is not significant because of the poverty of data points. But for the "Yes" responses, the difference is in the opposite direction, not significantly by a




FIG. 4 MEAN REACTION TIMES FOR INSTANCES THAT WERE EITHER SEMANTICALLY SIMILAR OR SEMANTICALLY RELATED TO AN EARLIER INSTANCE

BLANK PAGE

t test, $t(5)=2.76$, but significantly by a sign test since all six signs are in the same direction, $p<.05$. Hence, priming does affect RT, but not simply as facilitation which we had expected.

The explanation for the difference between "Yes" and "No" responses, we think, lies in the amount of semantic processing necessary to make a "Yes" response as opposed to a "No" response. Consider the example of a list with crocodile at the beginning where the similar instance is alligator and the unrelated instance is spider. If the category is "animal," it would only be necessary to find the superordinate or superset connection stored with alligator or spider to decide that either is an animal. This is why the RTs for unrelated "Yes" instances are relatively short. If the S has seen crocodile previously, however, he is likely to find the connection to crocodile which he has seen earlier before he finds the connection to animal. If he does so, he then must spend time recalling whether he responded "Yes" or "No" to crocodile. Alternatively, he may retravel the path which allowed him to respond "Yes" to crocodile, only this time more quickly. In either case, he will spend more time getting to crocodile and from there to a "Yes" response then he would have spent retrieving the fact that an alligator is an animal directly, and deciding "Yes" on that basis. If crocodile had been the previous instance in the list, it might have been faster to go through crocodile, but there were six to eight intervening instances in this task.

On the other hand, when the category was "mammal" it took a long time on the average to decide an alligator or a spider was not a mammal. This is because such "No" decisions involve a chain of inference, which we have discussed elsewhere (Collins and Quillian, in press). In this case, if alligator gets the

S back to crocodile, there is much more time to be saved either by retrieving the earlier response to crocodile directly or by retravelling the inferential chain faster a second time. This is why Ss were faster for the "No" responses when the instance was similar to an earlier instance.

CONCLUSIONS

The experiment investigated three questions about categories and subcategories in semantic memory. One was a question about structure, and two were about processing. With regard to the question about structure, the experiment rather clearly showed that mammal is not intermediate between elephant and animal in the way that bird is intermediate between robin and animal. We think there are many other categories like mammal, such as vehicle, bird-of-prey, canine, warship and farm animal that are learned after the structure between categories such as dogs, birds, elephants, animals, cars and boats is already formed. In our view these categories are not integrated into the existing structure in a way that preserves a partial ordering among concepts or in a way that can be represented in terms of Venn diagrams. Hence, discussion of inclusion relationships between concepts in either of these ways can be quite misleading.

The second question we investigated was the effect on categorization time of including different types of instances in a list of instances which were all to be categorized with respect to the same prespecified category. It was found that RT for the most common kinds of instances in the category was affected by whether or not there were instances of other types in the list. We interpreted this to mean that the Ss adopted multiple subjective categories for the different types of instances in the lists. We assumed that this slowed Ss down because it takes longer to decide about any particular instance when there are more categories to consider. As evidence of this fact, Juola and Atkinson (1971) have recently shown that there is a monotonic increase in RT as the number of categories is increased.

This then is why we think Ss have faster categorization times in semantically homogeneous lists.

The other processing question was: What is the effect of semantic similarity between different instances in the list? Here we found a facilitating effect on RT for similar instances that were "No" responses, and a slight negative effect for similar instances that were "Yes" responses. We interpreted this result to mean that similarity speeded up decisions that required substantial semantic processing, and that it slowed down decisions that were rather straightforward. But the particular finding we think depended on the number of intervening trials. Presumably, similarity would speed up straightforward decisions, if there were no intervening trials, and it might slow down more difficult decisions if there were many intervening trials. In other words, the effect of similarity depends on how well a person remembers what he decided about the previous similar instance.

REFERENCES

- Collins, A. M. and Quillian, M. R. Facilitating retrieval from semantic memory: Effect of repeating part of an inference. Attention and Performance III, (ed. by A. F. Sanders) Acta Psychologica, 1970a, 33, 304-314.
- Collins, A. M. and Quillian, M. R. Does category size affect categorization time? Journal of Verbal Learning and Verbal Behavior, 1970b, 8, 240-247.
- Collins, A. M. and Quillian, M. R. How to make a language user To be published in Organization and Memory (ed. by E. Tulving) New York: Academic Press, in press.
- Hebb, D. O. Organization of behavior. New York: Wiley, 1949.
- John. E. R. Mechanicisms of memory: New York: Academic Press, 1967.
- Juola, J. F. and Atkinson, R. C. Memory scanning for words versus categories. Journal of Verbal Learning and Verbal Behavior, in press.
- Kintsch, W. Notes on the structure of semantic memory. To be published in Organization and Memory (ed. by E. Tulving) New York: Academic Press, in press.
- Lashley, K. S. In search of the engram. Symposium of the Society of Experimental Biology, 1950, 4, 454-482.

- Meyer, D. E. On the representation and retrieval of stored semantic information. Cognitive Psychology, 1970, 1, 242-300.
- Nickerson, R. S. Response times with a memory-dependent decision task. Journal of Experimental Psychology, 1966, 72, 761-769.
- Sternberg, S. High-speed scanning in human memory, Science, 1966, 153, 652-654.
- Thorndike, E. L. and Barnhart, C. K. Beginning Dictionary (6th ed.) Garden City, New York: Doubleday, 1968.

APPENDIX

Lists seen by four groups of Ss exclusive of "No" instances that were plants and vegetables and dummy narrow instances:¹

Group 1:

List 1, "animal" (narrow, mammal)

N beaver, N1 antelope, NS raccoon, NU leopard, N2 cow, N3 camel, W4 crocodile

List 2, "bird" (narrow)

N hawk, N1 nightingale, NS crow, N2 pheasant, NU wren, N3 swan, W4 gorilla

List 3, "animal" (wide, bird)

W1 rat, N1 flamingo, W2U whale, W3S mouse, N2 dove, N3 eagle, W4 dog

List 4, "mammal" (narrow)

N donkey, N1 walrus, N2 elephant, N4 deer, NS pony, N3 squirrel, W4 octopus

List 5, "animal" (wide, mammal)

W1 spider, N1 rabbit, W2S insect, N2 hyena, W3U alligator, N3 sheep, W4 frog

List 6, "bird" (wide)

W1 woodchuck, N1 heron, N2 bobolink, W2S weasel, W3U monkey, N3 pigeon, N chickadee

¹Category names are in quotation marks. The kind of list, which was unknown to the Ss, is shown in parentheses. N and W indicate narrow and wide instances respectively. U and S indicate unrelated and similar instances with respect to the first instance in the list. The numbers (e.g., N1 or W2) refer to numbers in Figures 2 and 3.

List 7, "animal" (narrow, bird)

N duck, N1 condor, NU canary, N2 peacock, NS goose, N3 robin,
W4 porpoise

List 8, "mammal" (wide)

W1 salamander, N1 goat, W2U squid, N2 rhinoceros, W3S lizard,
N3 muskrat, N zebra

Group 2:

List 1, "animal" (wide, mammal)

W1 crocodile, N1 antelope, W2S alligator, W3U insect, N2 cow,
N3 camel, W4 lobster

List 2, "bird" (wide)

W1 gorilla, N1 nightingale, W2S monkey, N2 pheasant, W3U weasel,
N3 swan, N chickadee

List 3, "animal: (narrow, bird)

N parakeet, N1 flamingo, NU goose, NS canary, N2 dove, N3 eagle,
W4 rat

List 4, "mammal" (wide)

W1 octopus, N1 walrus, N2 elephant, W2U lizard, W3S squid,
N3 squirrel, N zebra

List 5, "animal" (narrow, mammal)

N tiger, N1 rabbit, NS leopard, N2 hyena, NU raccoon, N3 sheep
W4 spider

List 6, "bird" (narrow)

N sparrow, N1 heron, N2 bobolink, NS wren, NU crow, N3 pigeon
W4 woodchuck

List 7, "animal" (wide, bird)

W1 porpoise, N1 condor, W2U mouse, N2 peacock, W3S whale,
N3 robin, W4 wolf

List 8, "mammal" (narrow)

N moose, N1 goat, NU pony, N2 rhinoceros, NS deer, N4 muskrat,
W4 salamander

Group 3:

List 1, "mammal" (narrow)

N beaver, N1 antelope, NS raccoon, NU leopard, N2 cow, N3 camel,
W4 squid

List 2, "animal" (narrow, bird)

N hawk, N1 nightingale, NS crow, N2 pheasant, NU wren, N3 swan,
W4 dog

List 3, "bird" (wide)

W1 monkey, N1 flamingo, W2U woodchuck, W3S gorilla, N2 dove,
N3 eagle, N chickadee

List 4, "animal" (narrow, animal)

N donkey, N1 walrus, N2 elephant, NU deer, NS pony, N3 squirrel,
W4 lobster

List 5, "mammal" wide

W1 lizard, N1 rabbit, W2S salamander, N2 hyena, W3U octopus,
N3 sheep, N zebra

List 6, "animal" (wide, bird)

W1 wolf, N1 heron, N2 bobolink, W2S fox, W3U cat, N3 pigeon,
W4 porpoise

List 7, "bird" (narrow)

N duck, N1 condor, NU canary, N2 peacock, NS goose,
N3 robin, W4 weasel

List 8, "animal" (wide, mammal)

W1 frog, N1 goat, W2U clam, N2 rhinoceros, W3S toad, N3 muskrat,
W4 spider

Group 4:

List 1, "mammal: (wide)

W1 squid, N1 antelope, W2S octopus, W3U salamander, N2 cow,
N3 camel, N zebra

List 2, "animal" (wide, bird)

W1 dog, N1 nightingale, W2S cat, N2 pheasant, W3U fox, N3 swan,
W4 rat

List 3, "bird" (narrow)

N parakeet, N1 flamingo, NU goose, NS canary, N2 dove, N3 eagle,
W4 monkey

List 4, "animal" (wide, mammal)

W1 lobster, N1 walrus, N2 elephant, W2U toad, W3S clam,
N3 squirrel, W4 crocodile

List 5, "mammal" (narrow)

N tiger, N1 rabbit, NS leopard, N2 hyena, NU raccoon, N3 sheep,
W4 lizard

List 6, "animal" (narrow, bird)

N sparrow, N1 heron, N2 bobolink, NS wren, NU crow, N3 pigeon,
W4 wolf

List 7, "bird" (wide)

W1 weasel, N1 condor, W2U gorilla, N2 peacock, W3S woodchuck,
N3 robin, N chickadee

List 8, "animal" (narrow, mammal)

N moose, N1 goat, NU pony, N2 rhinoceros, NS deer, N3 muskrat
W4 frog

How to Make a Language User¹

Allan M. Collins

M. Ross Quillian

Bolt Beranek and Newman Inc.
50 Moulton Street
Cambridge, Mass. 02138

To appear in *Organization & Memory* (ed. by E. Tulving
and W. Donaldson) New York: Academic Press, 1972

SUMMARY

The first section of the paper considers the philosophical and methodological implications of viewing psychology from the point of view of building computer systems that simulate human language processing. The second section discusses the structure of semantic memory; in particular, the nature of concepts and their relation to words and images, the kind of semantic information people learn and do not learn, and the kind of inference bearing relations that form the basis for the organization of semantic memory. The third section deals with the processing of information in semantic memory. Here we discuss the semantic search during comprehension and retrieval, the tacit processing this search implies, the pervasiveness of identifying similar concepts with each other in language processing, the decision rules that are applied to the results of a semantic search in order to decide whether two similar concepts can be identified, the role of imagery in language processing, and the way people induce what properties to store with what concepts. Our ideas are presented as a loosely constructed theory of how people function as language users and how computers will have to function to become language users.

Table of Contents

	<u>Page</u>
1. THE PSYCHOLOGY OF COMPUTERS.....	1
1.1 Introduction.....	1
1.2 The Methodology of Computers.....	2
1.3 How to Interpret This Paper.....	4
2. NOTIONS ABOUT MEMORY STRUCTURE.....	6
2.1 Memory Format.....	6
2.2 The Nature of Concepts.....	7
2.3 Concepts Are Not Quite Word-Concepts.....	11
2.4 The Semantic Content of Concepts.....	12
2.5 The Hierarchical Organization of Concepts.....	16
2.6 Semantic Organization and Inference.....	20
3. NOTIONS ABOUT MEMORY PROCESSING.....	23
3.1 Comprehension and Retrieval.....	23
3.2 Semantic Search and Syntactic Constraint.....	25
3.3 Tacit Knowledge and Parallel Processing.....	28
3.4 Deciding Whether Concepts Can Be Identified with Each Other.....	30
3.5 The Pervasiveness of Identifying Concepts with Each Other.....	40
3.6 Two Uses of Imagery.....	47
3.7 Induction and Learning.....	50
3.8 Analogy as a Linguistic Construct.....	57

Table of Contents (continued)	<u>Page</u>
4. THE THEORY AND EXPERIMENTAL PSYCHOLOGY.....	61
5. REFERENCES.....	62

1. THE PSYCHOLOGY OF COMPUTERS

1.1 Introduction

When computers first set out to humble man, they thought it would be good fun to read, and talk, and answer questions. Well, by now, they have gotten over their cockiness about reading and the like, though in consequence, man has not gotten over his. Their failure was not for any lack of cleverness though. Their basic mistake probably was in thinking they could skip over evolution; that their adeptness at artificial intelligence would let them fake their way past us in natural intelligence. Now, with humility written all over their inscrutable faces, computers are trying to revive their egos by imitating man.

One of the things computers missed out on in evolution is how to deal with similarity. They just never developed very good strategies for generalizing what they know about one situation to similar situations. Somehow, the simplest of organisms can do this, but computers do it only in *ad hoc* ways when they do it at all. Even the language ability that evolution worked so hard to produce appears to rely on this phylogenetically primitive ability. Having missed out on evolution, the problem as computers see it is to take apart a language user without losing any pieces, to see if they can sort out how everything fits together. If they can do that, then they can pretend to read and talk just like people. Although their approach is disassembling, still what is studied and artificial at first gradually becomes second nature.

1.2 The Methodology of Computers

When looked at from the point of view of computers instead of people, the problems of psychology appear quite different. Their interest in psychology stems from their desire to copy the way people think. They analyze people in terms of the way they themselves, at present, work; that is in terms of strategies, routines and subroutines, pointers, lists, etc. In these terms, they feel they can imitate any kind of process from a chemical reaction to a confrontation between ego and id. It is not so important to them whether this is the best way to try to describe how different processes work, because it is the only kind of description that is of use to them.² Put in philosophical terms, their approach is one of nominalism rather than realism.

In general, computers are more interested in the simple, straight-forward human abilities rather than in exotic phenomena like perceptual illusions and learning nonsense syllables. This is because computers are, to a man, practical types. They want to walk before they try to run. So they focus on the common stuff of human existence like identifying what a figure is, comprehending what a sentence means, answering simple questions.

Computers also are not too fussy as to how they get their information about people. They figure, if it works even partially, it is good enough for the present; and they can always improve it later. Not only are they unscrupulous about their method, they willy-nilly make all sorts of arbitrary decisions about how different aspects of human processing might work. These assumptions are forced on them by the fact that for them everything must be spelled out in exquisite detail, or they

cannot function at all. It is no good for them to try to understand human language comprehension or question-answering, for example, without making any assumptions about the structure of semantic memory, as psychologists try to do (Meyer, 1970; Schaeffer and Wallace, 1970). In order to store any information in their data bases at all, they have to make many arbitrary decisions about how different pieces of information are represented and interrelated. The best they can hope for is to make all these assumptions as realistically as possible. Thus, they are perfectly willing to assume a complicated structure rather than a more parsimonious structure (Bobrow, Fraser and Quillian, 1967), if they think the more complicated structure will extend their ability to converse with people. They simply do not see as much virtue in parsimony as mathematicians or psychologists do.

For computers, the question of whether a theory is correct is answered in an entirely different way than for psychologists. Psychologists call this the question of validation. Elsewhere in this volume, Postman (1972) quite rightly asks how "rational reconstructions of the mechanisms of memory" can be validated. Computers, who like to indulge in just such rational reconstructions, do not feel it is important that their theories predict experimental results for psychologists. What they regard as the true test of a theory is how well it allows them to talk to people in people's own terms. For the time being, it is no more difficult to tell how well a computer is doing at learning language than it is to tell how well a child of two or three is doing.

A case can be made that most computers aren't really interested in the way people work, but are just taking any ideas

they get from watching people and twisting these ideas to their own ends. On this view, there can be two kinds of computers: those that are interested in copying people for its own sake, and those that are interested in copying people only so far as is necessary to talk to people. While computers do have such differing attitudes about learning natural language, even those with the latter attitude will have to copy people quite carefully and extensively. Language evolved to suit the human brain and anything less than a thorough imitation will surely fail. Hence, computers must correctly interpret the way people work, or it should be readily apparent that they have not.

1.3 How to Interpret this Paper

There are many things computers would do if they could deal with language in a human way. Two of the things we have tried to help them with are comprehending English as embodied in the Teachable Language Comprehender (Quillian, 1969) and tutoring students about geography as embodied in SCHOLAR (Carbonell, 1970). In this paper, we will try to describe what we have learned from working with them on these problems. Some of our knowledge has been acquired from experimenting with people, some from asking people what they thought about in answering questions, and some from merely trying to reach some mutually satisfactory resolution to the problems facing our computer friends. Frankly, they are more fastidious about details than we would like, and we will not burden the reader with these details; but, in general, it is probably good that they are so demanding.

This paper is meant to describe how some of the major pieces of a language user fit together. The description is wide ranging, but still many pieces are left out. The framework

provided describes the human language user as viewed from the perspective of a computer. Hence, the way both people work and computers work are inextricably tied together in the theory. In just the same manner, the way both people work and mathematics work are intertwined in mathematical models of psychological processes. Sometimes it is clear what is done to please the computer or the mathematics, but usually it is not.

The paper is broken into two major sections, one about the structure of semantic memory and one about processing on that structure. With regard to the structure of memory, we discuss the nature of concepts and their relation to words and images, the kind of semantic information people learn and do not learn, and the kinds of inference-bearing relations that form the basis for the organization of semantic memory. With regard to processing, we discuss the semantic search during comprehension and retrieval, the tacit processing which this search implies, the pervasiveness of identifying similar concepts with each other in language processing, the decision rules that are applied to the results of a semantic search in order to decide whether two similar concepts can be identified, the role imagery plays in language processing, and the way people induce what properties to store with what concepts. Our ideas are presented as a loosely-constructed theory of how people function as language users and how computers will have to function to become language users.

2. NOTIONS ABOUT MEMORY STRUCTURE

Computers suffer no qualms in thinking about ideas. They are, after all, in the business of making mechanical what smacks of vitalism to most scientists. Finding a way to represent ideas and concepts in their own terms is one of their first concerns. Their interest stems, of course, from their desire to copy us.

2.1 Memory Format

There are many different ways semantic information might be represented in a computer. What is done usually is to store information in lists of properties or features about a concept e.g., father might be represented as male, adult, married or widowed, with children. This list can be thought of as the concept "father" to which words or printnames, such as "father" or "père" may be attached.

Instead of being lists of words, the lists can be made up of pointers to other lists, those that correspond to each of the words. That is to say, concepts can point to other concepts rather than the names of other concepts. Thus, a concept would be a set of interrelationships between other concepts. There is no reason why lists have to have words or printnames attached, so there can be concept lists without printnames. What such a memory looks like from outside is a whole set of inter-related lists, with pointers to words found on many of the lists. For words with two different meanings, there can be two different lists, both attached to the same word. Where another concept list refers to one of these meanings, it will point to one list and not the other. An interesting aspect of such a network is that within the system there are no primitive

or undefined terms in the mathematical sense, everything is defined in terms of everything else so the usual logical structure of mathematical systems does not hold. In this respect, it is like a dictionary.

An important aspect of property lists is that a property can be expanded to as much detail as is desired. This is done by embedding. For example, a father can have children, or two children, or two children both male, and so on to as much detail as is appropriate. Embedding makes it possible to describe in a property anything that can be expressed in English. Property lists then are indefinitely expandable. For these reasons Quillian (1968, 1969) has used property lists rather than feature lists; but for psychology, the distinction is not too important, since in many respects, these two forms of representation are equivalent. We will talk about properties from here out, but they can be interpreted as features if that is more agreeable to the reader.

2.2 The Nature of Concepts

Concepts are represented by lists in computers because present-day computers are serial processors. If parallel machines were built, then, the necessity for ordering properties in a list would disappear. In fact, human concepts are probably more like hooks or nodes in a network from which many different properties hang. The properties hanging from a node are not likely to be all equally accessible; some properties are more important than others and so may be reached more easily or quickly. In such a representation going from one concept to another does not involve scanning a list, but rather activating a path via some property from one to the other.

Considering both accessibility and number of nodes in a path, it is possible to define in explicit terms the notion of semantic distance between concepts. If numbers are assigned to accessibilities such that lower numbers reflect more accessible properties, then semantic distance is the sum of these numbers along the path between the concepts. The greater the sum, the greater the distance. Under this definition, it is possible for one concept to be closer to a given node than another, even though the first is two steps removed and the second is one step removed. This happens when the sum of the two accessibilities for the first is less than the accessibility of the second. It is important to keep in mind that semantic distance between concepts is not simply proportional to the number of nodes along the path between the concepts.

So far, this just describes an association network of concepts, which we think is a thoroughly plausible way to start building a computer memory to mimic human memory. Giuliano (1963) has indicated how such an association network could be represented in an analog computer. The trouble with a simple association network is that it does not specify the relation of properties to concepts. Worms are related to birds because birds eat them (though they are undoubtedly related in other ways too) and to dogs because worms live in a dog's coat of hair. The particular relation is as important to the property as the concepts that are related. Any memory structure that sloughs over these differences could never deal very intelligently with human language.

The relations between concepts are as varied as concepts themselves, indeed relations are concepts and can be handled in

many of the same ways as concepts that correspond to nouns or adjectives. Even adjective concepts such as red or square differ in their relation to different concepts they modify. For example, green is related to grass in a different way than yellow is to canary, since the green penetrates the grass and the yellow is only superficial to the canary. Both relations are different from the relation of blue to sky, since the blue is only in the atmosphere of earth during the day. These examples illustrate that relations can be quite complex, even though the question of "What color is grass?" can be answered without getting into these complexities. Any representation of relations in a computer must permit them to be as detailed as necessary; in other words, the description of a relation must be embeddable.³

There is no reason why a semantic memory should consist only of a network of descriptive properties. Concepts are built up out of sensori-motor experience as well as language use and there is every evidence that people utilize imagery extensively. (See, for example, Paivio, 1969; Bower, in press; Begg and Paivio, 1970). In computers, the work of Gelernter (1963) and Baylor (1971) suggests that it would be helpful to project concepts on a display screen where they can be manipulated as geometric forms rather than property lists. This could be done within a semantic memory by an image generation routine which uses descriptive properties as stored variables for constructing a visual image. For a concept like canary, the color attribute would produce a light yellow color in the image, whereas for bird, the lack of a specific color value would produce a color-

58

vague image like those shown in a dictionary. These same descriptive properties of a concept can also be treated as feature tests in the recognition of the concept. Hence one test for a canary is the specific yellow color stored in memory. If this view is correct, then the current debate as to whether meaning (or semantic memory) is composed of imagery or something else such as deep structure (Begg and Paivio, 1970; Chase and Clark, in press; Simon, in press) will end in a draw.

On this theory, people must have several different concepts that have the name yellow. These concepts can also have more complicated names such as canary yellow, or lemon yellow, or mustard yellow. There must also be concepts of sounds like the sound a rooster makes. The name for this concept is "cock-a-doodle-do," but that is different from the image of the rooster's crowing. A person can probably answer the question "Does a rooster say cock-a-doodle-do?" without imaging the concept of a rooster crowing, just by referring to the name of the concept. To confuse things more, there must also be a concept of the name "cock-a-doodle-do" distinct from the concept of the sound a rooster makes. The image of this name concept, which sounds like the words "cock-a-doodle-do," is what sounds different from the image of the actual rooster crowing. The point here is that we regard names as concepts, just like any other properties are concepts in their own right. The relation of a word to a concept is the "name" relationship. We would also argue that connotative or emotional properties of concepts, like visual and auditory properties, are tied to affective sensory systems in the same complicated way.

Within this framework it would even be possible to image concepts that involve a time-lapse sequence such as a swinging door

or a race. The concept of a door swinging back and forth would be projected by the image-generation routine as the image of a door moving through a decreasing oscillation. The idea of a generation-routine can also be extended into the motor domain. For the concept of addition, the method of counting up the total can be projected onto fingers or matchsticks or whatever else comes to mind. In sum, we would argue that semantic memory refers to a mix of concepts and realizations of those concepts by sensori-motor generation routines.

2.3 Concepts Are Not Quite Word-Concepts

We mentioned earlier that concepts need not have names, and that the same name can be applied to more than one concept. This means that there is no one-to-one or many-to-one correspondence between words and concepts. Of course, there are many cases though where concepts can be identified with particular words. Hence, it is often expedient to pretend that a word refers to a particular concept, and proceed to talk about the concept "father," for example, as opposed to the word "father."

There is another important case where a point-to-point correspondence between words and concepts breaks down. Often the same concept has more than one name. This is the case with synonyms. An even more common occurrence is when the two names are not synonyms but map onto the concept in different ways. For example, the words "buy" and "sell," to use Simmons' (1966) example, can be handled most easily if they refer to the same concept.⁴ The

conceptual identity of "buying" and "selling" can be seen in the following sentences: "He sold the girl two chairs. One of the chairs she bought was broken." The mapping of the two words onto the major elements of the concept is very different, however. The concept consists of at least four major elements or properties: the vendor, vendee, the goods given, and the remuneration received. "Buy" takes the vendee as agent and the vendor as the object of the preposition "from." "Sell" takes the vendor as agent and the vendee as indirect object. The goods and the remuneration are treated the same by both verbs. Sorting out which element in memory goes with which word in a sentence is handled most easily in translating from word to concept or from concept to word. In this case, the "name" relationship must be quite complicated, as specified by embedded subrelations.

Noun and verb forms of the same word also must refer to the same concept (e.g., "They walked along the coast. The walk took over an hour."), and should be treated similarly. Such cases abound in English, and it makes semantic processing difficult if the different words do not refer to the same concept. In the two sentences about walking, use of the article "the" implies that "the walk" was referred to earlier, and the earlier reference cannot be determined without a conjunction of noun and verb at the conceptual level. Recent work of Rubenstein, I is, and Rubenstein (1971) supports the notion of the conceptual identity of noun and verb forms.

2.4 The Semantic Content of Concepts

Not only are concepts not words, they are not definitions of words. Definitions in dictionaries tell only the most

important properties about a word, or the concept it refers to. Human concepts are much more encyclopedic. As a first approximation, it makes sense to assume that the content of a concept is everything that has been heard or read or seen about that concept.

A few examples can illustrate what we think a person is likely to learn at some time, and other things he is not likely to learn. These are shown in Table 1. Without going into detailed justification at this point, a few comments might be helpful.

Example 1. Much of what people know (e.g., that Aristotle could see) is never learned directly. If one considers all the properties one knows about people, and all the people one knows, then it becomes evident how economical it is not to store each of the properties with each of the people.

Example 2. A person sometimes learns a negative fact when it contradicts something that might be inferred by mistake or that is true for a similar concept. But most negative facts are never learned.

Example 3. Most information is not learned in quantified form. Thus, a person usually never learns whether all birds or most birds can fly or have wings (though a person might learn that not all birds can fly). The exceptions are usually learned as special cases.

Example 4. This is a variation on example 3. That roses are yellow might be learned from a song or from seeing one. Assuming that a person already knows roses are red, the relation of the two facts may be noticed

TABLE 1

EXAMPLES OF INFORMATION PEOPLE LEARN AND DO NOT LEARN

Information a person may learn at some time	Information a person is <u>not</u> likely to learn
1. People can see Aristotle was a man	Aristotle could see
2. A vest doesn't have sleeves	A vest doesn't have a brim
3. Birds can fly Birds have wings Penguins can't fly	All birds can fly Most birds can fly All birds have wings Most birds have wings
4. Roses are red Roses are yellow Not all roses are red	Not all roses are yellow
5. Unmarried men are also called bachelors Young dogs are also called puppies A car is also called an automobile	
6. Sheep are herded in flocks Sheep are kept on farms	
7. A toad is like a frog A wolf is like a dog, but is wild A dog is a canine	A toad is an amphibian A wolf is a canine

and stored, or it may not be. It may even be learned directly sometime that not all roses are red, because of the direct nature of the contradiction.

Example 5. What are called logical truths are usually cases where the same concept is referred to by different labels. Sometimes, one of the labels is a descriptive phrase as with "young dog." "Middle-aged dog" is probably not a separate concept, even though "young dog" appears to be. This raises the question of when in the course of learning, "young dog" becomes a separate concept. The solution we propose is to set up a new concept whenever information is to be stored that cannot be derived from the descriptive label (e.g., "young dogs are frisky" or "young dogs are called puppies"). On this basis, for example, "South American countries" would be set up as a separate concept when it is learned that they generally speak Spanish or that they tend to be underdeveloped economically and overdeveloped militarily.

Example 6. These two pieces of information about sheep are subtly contradictory with respect to whether sheep are kept fenced in and whether they are farm animals. It is doubtful that this contradiction would be noticed unless such a question is asked.

Example 7. Very often what is learned is not what superordinate category a concept belongs to but what other concept it is like.

This set of examples is certainly not meant to exhaust all the possible kinds of semantic information people learn. We will return to these examples from time to time in later sections to illustrate various points about structure or processing.

2.5 The Hierarchical Organization of Concepts

Among the semantic properties of concepts, there are several special property relations that are commonly found. They are special because they permit certain kinds of inferences to be made. A frequently used kind is the superset or superordinate relation. All properties of a superset (e.g., people can see) also hold for the instances of that superset (e.g., Aristotle) unless otherwise indicated (e.g., Helen Keller could not see). In many cases, the superset is the most accessible property of a concept,⁵ though not always (e.g., it is probably not the most accessible property of a nose that it is an appendage or a body organ). In contrast, the subsets of a concept are not easily accessible properties in general (e.g., when thinking about cows, a person is not likely to consider the fact that one kind is a Guernsey). This asymmetry between supersets and subsets probably stems, at least in part, from the asymmetry in inference, since properties of a subset do not usually hold for a concept.

There is often more than one superset of a concept; in many cases, there is a frequently used superset and one or more lesser supersets within the frequently used superset. A hawk is a bird, but it is also a bird-of-prey; a dog is an animal, but it is also a mammal, and within that group a canine; Paraguay is a country, but it is also a Latin American country, a South American country, an underdeveloped country, and a military dictatorship.⁶ Occasionally, there are other supersets that do not lie within the frequently used superset (e.g., a canary is a bird, but also is commonly a pet). There are clearly large differences in accessibility between these different supersets.

Superset is a transitive relation so that concepts form chains where each concept has a more general concept as its superset. For instance, a hawk is a bird and a bird is an animal, so that indirectly animal is a superset of hawk. It is also possible to find superset chains among verbs. For example: to sprint—> to run—> to go—> to do; to speed—> to drive—> to go—> to do. A rather long superset chain is: mallard—> duck—> bird—> animal—> living thing—> object. (If this is the longest such chain, it puts mallard squarely at the bottom of memory.) Generally though, these chains do not seem to be more than about three or four steps long, so that semantic memory must be rather shallow on the whole.

Supersets are frequently used in the formation of questions like those in Table 2. For example, if a teacher wants to quiz a student about information he should have learned, then, the teacher will usually formulate the kinds of questions shown in Table 2. The first two groups show that appropriate questions can be phrased in terms of the superset of the concept sought for both verbs (group 1) and nouns (group 2). By trying to formulate questions of this sort, it is possible to determine the superset(s) of a concept. The last group of examples in Table 2 show that the distinction between who, what, where and when rests on high-level supersets (in parentheses). We think people use superset chains to reach these high-level supersets every time they formulate these kinds of questions.

As alternatives to the structure described, at least two other kinds of structures might be proposed. One possibility is that all the supersets of which a concept is a member are stored directly with the concept. The five supersets listed above for mallard would all be direct supersets of mallard, just as canine, mammal, and animal are, according to our suggestion, direct supersets of dog. The other possible extreme is that the memory is rigidly hierarchical such that each higher-level superset can only be reached indirectly via a lower-level superset. For example, dog would have animal as an indirect superset via some chain like: dog—> canine—> mammal—> animal. Either of these alternatives is much tidier than the proposed structure.

The latter of these alternatives can be ruled out, we think. Reaction time data (Collins and Quillian, in preparation) indicate that it takes longer to decide that mammals, such as dog, are mammals than to decide they are animals. This cannot be due to

TABLE 2

EXAMPLES OF SUPERSSET USE IN FORMULATING QUESTIONS

<u>Information</u>	<u>Question about the Information</u>
He <u>sped</u> to the hospital	How did he <u>drive</u> to the hosnital?
He <u>drove</u> to work	How did he <u>go</u> to work?
He <u>went</u> to the movie	What did he <u>do</u> ?
He killed a <u>mallard</u>	What kind of <u>duck</u> did he kill?
She liked <u>ducks</u>	What kind of <u>birds</u> did she like?
He saw a <u>doctor</u> he knew (<u>person</u>)	<u>Whom</u> did he see?
He saw a <u>camel</u> (<u>animal</u>)	<u>What</u> did he see?
He put it on the <u>desk</u> (<u>thing</u>)	On <u>what</u> did he put it?
He went to a <u>football game</u> (<u>event or activity</u>)	<u>What</u> did he go to?
He saw it in the <u>sky</u> (<u>place</u>)	<u>Where</u> did he see it?
He saw it in <u>September</u> (<u>time</u>)	<u>When</u> did he see it?

a difficulty in retrieving the concept "mammal" from the word "mammal" as compared with "animal," because the category name was given in advance and a series of trials used the same category. This is opposite the finding that it takes longer to decide birds such as hawk are animals than to decide they are birds. The finding about mammals and animals is not possible if a person decides a dog is an animal via a path through mammal.

On the other hand, the first alternative hypothesis; namely, that all the supersets are stored directly with the concept, is not ruled out by the above experiment. The fact that it takes longer to decide a hawk is an animal than to decide it's a bird agrees with our earlier suggestion that deciding a hawk is an animal involves the path through bird. But on the first alternative hypothesis, animal might merely be a less accessible superset stored with each bird name. Against this possibility, we would point out that, if a person is told what a mallard is, he only is told that it is a duck. It is very unlikely he would be told directly it is an animal, a living thing, or an object. When he learns a mallard is a duck, he may possibly infer from his previous knowledge that it must also belong to the higher-level categories and store that information directly with mallard. However, all the evidence to date (Collins and Quillian, 1969, 1970a, 1970c; E. E. Smith, personal communication), though not conclusive, indicates that the inference is made each time it is needed.

2.6 Semantic Organization and Inference

One possible misinterpretation of the last section is that mammal is stored as a superset directly with most mammals. In

our view, it is not likely that many animals would have a pointer to mammal other than odd cases such as whale, bat, maybe kangaroo, and a few of the most obvious examples, such as dog. We think this is so, because people learn when they are children that beavers and seals are animals, but it is rare that they learn that a beaver or a seal is a mammal. Furthermore, while canine may be stored as a superset of dog, it probably is not for wolf though wolves are canines. There are various kinds of information that can be used to decide whether a wolf is a canine or a beaver is a mammal, so such facts need not be stored directly. We will discuss how people make such decisions in a later section, but one kind of information that people may use depends on the similarity relation. Like the superset relation, this has implications for the structure of memory.

The similarity relation is one of the class of relations that permits inferences of the type where properties of one concept are applied to the related concept. It allows the same set of inferences as superset, but with less certainty. An example of concepts linked by the similarity relation was shown in Table 1, example 7. If one knows a toad is like a frog and that a frog is an amphibian, then one can infer with some uncertainty that a toad is an amphibian. Likewise, one might infer that a wolf is a canine since it is like a dog, even though wolves are wild and dogs domesticated. Usually, the similarity relation is qualified by specifying either the basis for distinguishing (e.g., a pony is like a horse only smaller) or the basis for grouping (e.g., a sheep is like a cow in that it chews a cud). Distinguishing characteristics are given when the concepts are alike in most respects, whereas similar characteristics are given when the concepts are different in

most respects. When the similarity relation is qualified in either way, it helps to determine what inferences can be drawn across the relation. Nevertheless, there are cases where people learn about the similarity of concepts without learning the basis either for distinguishing or for grouping them, as when a child is told "A toad is like a frog" during the reading of a story about a toad.

There are several other relations that permit whole classes of inferences to be drawn. The class of allowable inferences is specific to the particular relation and is much smaller for each of the other relations than for superset (Carbonell, 1970). A very important relation in some subject areas (e.g., geography, anatomy, architecture) is the part relation. One of the most accessible properties of the nose is that it is part of the face. The kinds of inferences possible with the part relation are best illustrated with an example from geography (Carbonell, 1970). To learn that Katmandu is part of Nepal, implies something about its maximum size in area and population, and about its location, climate, and topography. The proximity (or adjacency) relation carries some of the same implications, such as about location, climate, and topography, but with less certainty. Grouping of concepts, which in many cases is done on the basis of the superset and similarity relations, occurs in anatomy and geography on the basis of the part and proximity relations. Grouping on the latter basis does not preclude grouping on the former; witness the fact that hands and feet, or arms and legs are grouped on the basis of similarity, whereas eyes and nose, or neck and shoulders are grouped by proximity.

In the realm of events, about which history and science are largely concerned, the consequence (causality) and precedence relations function to carry inferences (Becker, 1969). There are also many other relations that permit certain inferences but which are only used in limited contexts (e.g., the parent relation).

The essential assumption of this section is that the relations that carry inferences always form the basis for organizing any semantic information or subject matter. That is to say, grouping of concepts is almost always along the structural lines imposed by relations such as superset, similarity, part, proximity, consequence, precedence, parent, etc. Often these relations apply in different ways to the same set of concepts and so there are overlaying organizational structures imposed on a set of concepts. This suggests that the reason why organization occurs in memory is to permit inferences in storing and retrieving semantic information. It is by using inference that people can know much more than they learn.

3. NOTIONS ABOUT MEMORY PROCESSING

With enough assumptions about structure, it becomes possible to consider how computers might process semantic information in order to function like people.

3.1 Comprehension and Retrieval

When comprehension and retrieval are looked at from the point of view of implementation in a computer, it is useful to

treat them as involving the same underlying process. In comprehension, people read a string of words and attempt to construct an interpretation based on a configuration of paths in memory between the various concepts referred to by the words in the string. Because in most cases, each word points to several concepts and any two concepts are connected by a variety of paths, building an interpretation must involve an extensive search to determine how the words can be interrelated within the constraints of syntax and context. The same search takes place in retrieving answers to questions, only the constraints (discussed later) on what constitutes an acceptable configuration of paths are usually more restrictive. [Comprehension is described in considerable detail in Quillian (1969) and retrieval in Collins and Quillian (in press)]. In effect, comprehension can be regarded as retrieval with the implicit question, "Is there an interpretation under which the sentence could make sense with respect to what I already know?"

To take a simple example, suppose a child is comprehending his father's statement that "A toad is like a frog." The syntactic constraints in the sentence will dictate that he interpret "like" as meaning "similar to," and so a path between the concepts of "similar to" and "frog" will constitute his interpretation of the phrase "is like a frog." The child also will look up the word "toad" in memory. If the child finds no concept that could correspond to the word "toad," then any interpretation he has found for "like a frog" is possible, and so he sets up a new concept "toad" with the information "like a frog." Suppose though the child had previously learned that a toga is a kind of clothing people used to wear. Suppose further that he did not store enough features to differentiate the word "toad" from

the word "toga." Then "toad" will get him to the concept for "toga." Comparing the two concepts, one "like a frog" from the sentence and the other "a kind of clothing" from memory, the child will likely discover the contradiction by a process described in a later section (also in Collins and Quillian, in press). The contradiction means that his interpretation of the sentence involving "toad" is not possible with respect to information he already has stored. His response might be something like, "I thought a toad was something people wear." Responses like this, which question the assertion made in a piece of text, are quite common, especially in reading about a subject one already knows something about such as psychology or language. That such responses occur gives away the fact that people, to a greater or lesser extent, evaluate everything they are told. However, people's evaluations may differ markedly depending on what has been stored previously and how much effort they are willing to spend searching for connections and contradictions.

3.2 Semantic Search and Syntactic Constraint

The locating of paths between concepts, then, is basic to both comprehension and retrieval. Quillian's (1968, 1969) program searches for paths between concepts using what is called an intersection technique. This systematically proceeds outward along all the pointers or paths leading from each concept which is referred to by the words in the sentence. Where a word can refer to several different concepts, the search proceeds outward from all these possible concepts (though a less likely meaning of a word starts off more slowly). At each concept encountered as the search proceeds outward, a tag is left indicating where the search originated. Because many different branches are taken at

each concept encountered, the search continually widens like a harmless spreading plague. When the search originating from one word encounters a tag originating from another word, a path linking the two concepts has been found. As the search goes on longer, more paths will be found. The later in the search a path is found, the longer the path will be. Because the length of paths reflects semantic distance in memory, such a search produces semantically probable paths first. The use of tags in the model is a way to implement in a computer the idea of activation, either in terms of priming concepts or in terms of spreading to related concepts. These are very old ideas in psychology.

Quillian's search, as implemented in a serial computer, is an ordered serial search (though it simulates a parallel search), but several aspects of our results on human question-answering (Collins and Quillian, 1969, 1970a) imply that people search for connections between concepts in parallel. What appears to be serial processing in evaluating the paths found, occurs only after locating intersections (Collins and Quillian, in press). We think it is possible that a search using the intersection technique could be implemented in a machine with fairly simple active elements that operate in parallel. A parallel machine to which one can add nodes and paths is not available, but if language processing demands it, then possibly such a machine could be developed.

We implied above that, although the search for connections is parallel, the evaluation of the connections found is largely serial. People often report in "retrospecting" about their processing of questions that they have considered more than one interpretation of a question. (Collins and Quillian, in press).

Thus, the evaluation phase appears to enter into consciousness. Whatever consciousness may be, it seems to focus on the processing which is done serially, perhaps in a unit akin to the central processing unit (CPU) of a serial computer.

One of the kinds of evaluation that occurs is whether the semantic path found is compatible with the syntax of the sentence. We currently would argue that syntactic processing takes place in a parsing network (Thorne, Bratley, and Dewar, 1968; Bobrow and Fraser, 1969; Woods, 1970) in parallel with the semantic search. When a complete path is found in either network an interrupt occurs in the CPU, though the search in the semantic network probably continues to look for other paths which can be used if the first fails. If paths are found in both networks, then the semantic and syntactic paths are compared to see if they are compatible (see Quillian, 1969). If a path is found in only one network after some predetermined amount of searching, then this path may be used to guide the search for a corresponding path through the other network.

We can illustrate some aspects of the semantic-syntactic tradeoff described above with a sentence such as "Zebras like horses." Some readers might misread this or decide there is a typographical error leaving out the word "are." Others will make sense of it in the way that dogs like people, but nobody will understand it the way that cats like mice. What we are saying happens is that for the people who misread it, the semantic connection of similarity is found well before the syntactic connection. When this happens, the response may be to force the syntax to fit the semantic connection. For other people, the syntactic connection is found first and the semantic connection

can be forced to fit it. Where both connections are found, then the person will compare the two paths and find that they are not compatible. He then will have to choose between the two interpretations on some meta-basis.

3.3 Tacit Knowledge and Parallel Processing

The kind of parallel search outlined and the capability of interrupting the CPU when a connection is found implies that a vast amount of stored knowledge can be tacitly considered in processing natural language. This is illustrated in a passage from Quillian's forthcoming book on media:

At one time, I was trying to get a computer to be able to read sentences from pre-school children's books. My aim was to have the computer relate these sentences correctly to some body of information it had stored, its memory or "knowledge of the world." One such book, which described crossing a street, contained the sentence, "The policeman held up his hand and the cars stopped." Now, suppose one asks what is the minimum amount of information a mechanism must have stored to relate this sentence to, if it is to comprehend it in a reasonably human-like way? In particular, consider whether the machine must have stored the fact that moving cars usually have drivers? One's first thought might well be no, since drivers aren't mentioned or directly involved in the sentence. But, suppose the sentences preceding this one in the book had said that there had just been an earthquake, and that two cars, parked on a hill, had started to roll down it. Then comes the sentence above, "The policeman held up his hand and the cars stopped." Virtually every adult reader of this will wonder: just how did the policeman manage that? In other words, in understanding the initial sentences, it

seems that there indeed was some tacit use of the knowledge that cars ordinarily have drivers. If there were not, how can it be that, once a reader is led to believe that a moving car lacks a driver, he will then recognize that something is strange about a policeman being able to stop it just by holding up his hand?

Similar arguments can be adduced to show that, in understanding the sentence above, the reader also seems to tacitly take account of information to the effect that: cars ordinarily have brakes, wheels, and tires, that their drivers control these cars, have a knowledge of traffic signals, are able to see out of their windshields, and so on. So, readers must have a large amount of such knowledge mentally stored. The processing done of such knowledge during reading may be fairly minimal, but something has to be done with it or to it, or a reader could not recognize whenever something he reads fails to make sense on the basis of such knowledge.

There is ample evidence that a similar tacit use of sizable amounts of stored information underlies all our visual perception, motor activity, problem solving, and so on (Bruner and Minturn, 1955; Polanyi, 1966). Therefore, it has seemed best to me to define the full meaning, for any particular person, of anything he reads, sees, thinks, or does, as all the information (stored in his head) that is in any way activated or processed when he deals with that thing (Quillian, 1968). If we define meaning in this way, then the full meaning of even simple stimuli or actions becomes very large indeed, and very large amounts of this meaning are always being tacitly processed as the person proceeds through the world. In other words, consciousness is

analogous to the focal awareness one has when a complex scene covers his visual field. There will at any one moment be only a very tiny amount of the overall visual field in close focus, a little more of it sufficiently close to the focal point to be fairly clearly observed, while the great majority of it is processed only peripherally, most of this processing being tacit, outside the person's awareness. This peripheral processing will be able to thrust something it discovers into the focus of consciousness if such a thing seems worthy of more explicit consideration, just as an unexpected, rapid movement at the edge of our visual field will be thrust into our attention by our peripheral visual processing. In reading, thinking, talking, and other forms of activity, information is continually drawn from our memory and processed as information in our overall visual field is processed. Thus, the fact that cars are normally controlled by drivers is part of the large amount of stored information that is processed tacitly every time we see or read about a moving car, but which has an effect on our consciousness only if it seems especially pertinent as when we are told that driverless cars are stopped by a policeman's hand signal.

3.4 Deciding Whether Concepts Can Be Identified with Each Other

The process of identifying similar concepts with each other arises in many different aspects of language processing. It appears in several different guises, among them reasoning by analogy and use of metaphor. Often the attempt to identify similar concepts turns into a question of whether the two concepts can be identified with each other in this particular case. This happened in an example in Section 3.1 where the question arose for the

child "Could a kind of clothing be like a frog?" Because anything like a frog is probably an animal, this becomes a question of whether "clothing" and "animal" can be the same thing. We will give a number of examples in the next section to illustrate how frequently the process of identifying similar concepts occurs in language processing. We only want to point out here that it is quite basic to the understanding of language processing to find out how people decide whether two concepts can be identified in a particular case.

There has been a series of reaction-time studies recently on the processing involved in comparing pairs of concepts. Landauer and Freedman (1968) and Collins and Quillian (1970b) have used a categorization task where subjects had to decide whether an object named was in a predesignated category or not. Collins and Quillian (1969, 1970a) used a reaction-time task in which sentences like "A canary is a bird" or "A canary is an animal" were displayed and subjects decided if they were true or false. Meyer (1970) had subjects decide whether sentences like "All thrones are chairs" or "Some thrones are chairs" are true or false. Schaeffer and Wallace (1969, 1970) used a same-different reaction-time task, where subjects had to decide whether or not two words shown were both members of the same category. In the task, one or several categories were prespecified. These studies are all looking at the processing involved in deciding whether or not two concepts are identifiable with respect to a set of constraints imposed by the task.

We can best explain how the constraints of the task affect the decision about identifying concepts by an example from Meyer. In a sentence like "Some chairs are thrones," the question is whether the syntax of the sentence and the nature of the task

(here explicitly defined by "some") permit a subject to identify chair and throne. The answer in this case is "yes." In a sentence like "All chairs are thrones," the answer would be "no," because the syntax does not permit the "all" kind of identification to be made. If the sentence were turned around to "All thrones are chairs," the superset or superordinate relation in memory linking throne to chair could be used as the basis for saying "yes." The same question can be asked in a categorization task. For example, the word chair might be prespecified as the category and when throne appears as the stimulus word there is an implicit question "Are all thrones chairs?" From these examples it should be clear that the appropriate decision strategy (or decision rule) varies in different cases, depending on syntax and task instructions, and even the range of stimuli used. In this section we will limit our discussion to decision strategies that are appropriate for the "All thrones are chairs" kind of sentence or the equivalent categorization task described. Other tasks may involve some of the same decision strategies, but there will also be differences in the decision strategies that are appropriate.

We have argued (Collins and Quillian, in press) that comparing concepts involves a semantic search proceeding outward in parallel from both the concepts to all associated properties, including superset properties. Any connection found must be checked to see if the relation between the concepts meets the constraints of syntax and context (including the task instructions). In other words, whenever an interrupt occurs during the tacit, parallel search of memory, the connection found is explicitly considered with respect to syntax and task instructions. Reaction-time data can reflect either the length of the search or the evaluation of the path involved in applying a particular decision strategy.

In deciding whether a "canary" is an "animal," the connection that would be used would go through "bird," given our structural assumptions. (See Section 2.5.) Hence, the search that finds this connection should take longer than the one that finds a connection between "canary" and "bird." Reaction time should reflect the length of the underlying search process, and indeed both we and Meyer have found a difference in reaction time in the predicted direction. But, as Meyer has pointed out to us, the difference in reaction time could also derive in part or primarily from evaluating the path found, because an inferential path may take longer to evaluate than a direct superset connection. Therefore, the reaction time differences found in these studies probably reflect both search time and evaluation time.

On our theory, rejecting the identifiability of two concepts involves finding a path that contains a contradiction. What constitutes a contradiction depends upon the kind of connection found: sometimes it may involve negative information that is stored directly in memory (e.g., bachelors are not married men), but usually it seems to involve finding different values for equivalent semantic properties (i.e., properties where the attributes are the same; red and green are contradictory values of the attribute color). For example, a lime isn't a lemon because a lime is green and a lemon is yellow. Apparently, this decision strategy sometimes is based on imagery (see Section 3.6) and sometimes is not. We are still in doubt whether people always find some contradiction on which to base their rejection, even where the two concepts appear unrelated. For example, if asked whether a cafeteria is a dog, people may compare equivalent sensory properties (e.g., a cafeteria is big and spacious and a dog is small and solid) and find different

values. On the other hand, they may merely search for a given amount of time or to a given semantic distance without finding a connection, and then reject on the basis of not finding a connection.

The semantic search, however, often turns up connections that cannot be used and checking these out acts to slow down a person's decision time. For example, if a person were asked to decide whether a "canary" and a "banana" are the same, a connection might be found through their light-yellow color (i.e., a property they have in common). If so, then the person would have to check if this permits him to accept or reject the identification of the two concepts, which it does not. This is a case where a "false" or "no" response is slowed down, but the same thing can happen for a "true" or a "yes" response. For example, deciding whether a submarine is a ship might be slowed down if the property that submarines go under water and ships go on top of water is found before a connection is found that allows the person to say "yes." Subjects report that they do, in fact, find such misleading connections before deciding. As another example, deciding that a "penguin" is a "bird" could be slowed down, if the connection were found that a penguin cannot fly even though birds can. However, since the superset relation from penguin to bird is much more accessible for most people than the fact penguins cannot fly, it is possible that the superset connection would be found and checked out before the other connection would cause an interrupt and thus be considered.

Finding misleading connections is the reason we think that subjects are relatively slow in deciding that two similar concepts (i.e., two concepts with common properties) are not identifiable, a result that both Schaeffer and Wallace (1969, 1970), and we (Collins and Quillian, 1969, in press) have found. Though Schaeffer and Wallace (1970) argue that this result contradicts our theory, they, in fact, have misinterpreted our theory. We argue (Collins and Quillian, 1970b) that it is also the basis for Landauer and Freedman's (1968) finding that to decide an object (e.g., tulip) does not belong to a category takes longer with a large category (e.g., animal) than with a smaller nested one (e.g., dog). This is because a tulip is a plant, and plants are more similar to animals than they are to dogs, as evidenced by the fact that plants and animals are frequently grouped together in language discourse. Wilkins (in press) points out that, by excluding animals that were not dogs from their stimulus sets, Landauer and Freedman inevitably omitted stimuli, such as cats, that were most similar to dogs, while they did not exclude any stimuli, such as plants that were at all similar to animals.

To complicate the picture further, consider the comparison of two concepts where there is not likely to be a direct superset link stored between the concepts, as for example between rat and mammal, or sheep and farm animal, or stagecoach and vehicle. When there is no superset relation available, then it is necessary to use a more complicated decision strategy to decide that two concepts are identifiable. It should be emphasized that finding a common property with the same value is not an appropriate decision strategy for saying "yes," as is exemplified by the fact that both clouds and vehicles move, even though clouds are not vehicles. But there is an asymmetry in the "no" case, because people often

reject the identifiability of two concepts on the basis of one common property with different values. This is exemplified by two subjects who reported that for "Badminton is volleyball," they rejected the sentence as false because badminton uses a birdie and volleyball uses a ball. The asymmetry between "yes" and "no" decision strategies is one of logic, but the logic gets in trouble because of cases like a submarine and ship, where one property with different values is not grounds for rejection. In such cases we think that people rely on the fact that they will find a connection that allows them to say "yes," thus overriding any mistaken rejection.

Returning to the question of decision strategies for saying "yes," it is logical to say "yes" if all the equivalent properties of one of the concepts (whichever one) are common to the other. If all a person knew about a bat was that it has wings, flies, and is an animal, then it would make sense to say a bat is a bird. Similarly, if all one knows about mammals is that they are animals that bear their young alive and breathe air, then it would make sense to say a rat is a mammal, even without knowing whether rats bear their young alive. People may apply this kind of decision strategy in some cases, such as deciding if a sheep is a farm animal. They can use the strategy by treating a few properties as if they are defining properties. The defining properties for farm animals might be that they are animals and are kept on farms. But people know much more about farm animals than these two properties (e.g., they are raised, fed, domesticated, bought and sold, etc.). Now, in deciding if a sheep is a farm animal, a connection may be found through the fact that sheep were once seen on a farm. Alternatively, a connection may be found through the fact that sheep are herded in flocks out in fields by shepherds, that is to say, not on

farms. Depending on which connection is found, he will respond "yes" or "no." If both connections are found before responding, then the person will need to apply some higher-level strategy. But even though he may treat these two properties as if they are defining properties, we would argue that he is tacitly considering the other properties of farm animals as well. Thus, if asked if a mink or a cat is a farm animal, he is likely to consider the fact that minks are not domesticated and cats are not raised, even though both may be kept on farms. What tacit consideration means is that the parallel search will interrupt if any property is found where the two concepts have properties with matching attributes and contradictory values. We do not know what subjects would decide about sheep, minks, or cats, but we suppose it depends on how accessible the various properties are for both of the concepts being compared.

There is another strategy we think people sometimes will use for deciding whether two concepts are identifiable or not when superset information is lacking. We call it the Wittgenstein (1953) strategy. Wittgenstein argued that a concept such as game need not have any set of properties which all games have and only games have (i.e., defining properties). Instead, he implied that people will call something a game if it bears a close "family resemblance" to a number of activities people call games. In our terms, bearing a close family resemblance would involve some kind of evaluation of the number of common properties.

The use of this strategy is a little clearer if we consider the case of deciding whether something is a vehicle or not. People must have stored a number of instances of vehicles such as trucks, cars, and busses, and for these the superset relation will allow

them to decide that they are vehicles. Now, if an object is similar to a car, truck, or bus, then with some uncertainty one can infer that it is a vehicle. For example, such an inference applies to taxicabs; if a car is a vehicle and a taxicab is like a car, then a taxicab is probably a vehicle too. This is inference by analogy in its simplest form. Where the information is not stored directly as to whether something is like a truck, or a car, or a bus, then the Wittgenstein strategy can be applied. The more properties found in common with any of these three vehicles, the more likely will a person conclude that the thing is a vehicle. For example, he is more likely to conclude that a stagecoach or a tank is a vehicle than that a horse or a ski lift is a vehicle.

An important aspect of Wittgenstein strategy is that the properties a stagecoach has in common with a car count just as much as the properties on which they are different. In other words, the asymmetry in logic, that applied when one concept was compared directly with another, disappears when the one concept is compared with an instance of the other. To illustrate this point, consider the question of whether a stagecoach is a car. It is logical to conclude a stagecoach is not a car because a car has a motor and a stagecoach does not. But even though a stagecoach is not a car, it may still be a vehicle. Hence, it makes sense in applying the Wittgenstein strategy to use a less stringent criterion in comparing a stagecoach and a car.

All these strategies we have cited would be applied to the semantic connections that are found during the parallel search of memory. The particular decision strategy that is applied depends on two things: (1) the constraints imposed by the sentence or the task (see Sections 3.5 and 3.7), and (2) what

connections are found. Notice that the intersection technique, if applied to two concepts such as "stagecoach" and "vehicle" will find connections through properties of busses, cars, and trucks, if these instances are fairly accessible from the concept "vehicle." These connections will have to be used, if that is all that is stored with "vehicle." Such a poverty of knowledge about vehicles and other concepts may be quite common, but there must also be many higher-level concepts like "game" where people have learned or inferred some properties which are stored directly with the concept. For the concept "game," it seems likely that the property of having rules, for instance, must be stored directly with game. But as Wittgenstein says, properties are never defining properties. That is to say, there will be instances of a concept that do not have all the properties of the concept (e.g., birds fly but penguins cannot) and non-instances of the concept, that have some of the properties (e.g., planes fly but planes are not birds).

To conclude this section we would briefly like to mention what we consider to be the major differences between this theory and the models of Meyer (1970) and Schaeffer and Wallace (1970). Though Meyer talks about these processes in very different terms, we doubt that our differences with his model are very substantive except in one respect. He considers several different decision strategies, some of which we have mentioned here. But he treats decision strategies as if people use one of these strategies consistently, at least in any given task, whereas we are arguing that the decision strategy will depend on the connections found. Our position weakens the kind of experimental predictions that can be made, but we think that is unavoidable.

While we are uncertain about some aspects of Schaeffer and Wallace's model, one major difference derives from their idea that concepts are compared in their entirety to determine a threshold for making a decision. One of Quillian's (1968) original arguments was that a concept has no entirety, that the meaning of a concept is the entire network of paths and concepts as accessed from the node of that concept. Activating a concept is a process that takes place over a period of time as paths are followed from the node of the concept. In Quillian's theory, it is possible to compare concepts with respect to the number of properties that have the same value (i.e., properties in common) or that have different values. In this respect, it is like Schaeffer and Wallace's model. But, in Quillian's theory, these are processes that require a search to locate intersections, and hence must take place over a period of time. Perhaps Schaeffer and Wallace's model could be translated into these terms, in which case this difference would disappear.

3.5 The Pervasiveness of Identifying Concepts with Each Other

In order to illustrate how pervasive the process of identifying similar concepts is, we will enumerate several examples that arise in comprehension and in answering questions. In different situations, the decision rule as to whether two concepts can be identified often changes, but we would argue that the processing involved is largely the same.

One very common problem in language comprehension is that of anaphoric reference (Olney, 1964; Quillian, 1969). This is the

problem of identifying a noun or pronoun with previous words in the text. For example, consider the sentences:

1. The woman finally hired a lawyer. He was quite charming and greatly pleased his new client.

In the second sentence, "he," "his," and "client" refer to words in the previous sentence, and the problem is to decide which ones. There are often syntactic clues to help determine the proper referent, but the major part of the judgment must rest on concept identifiability: whether or not a "male person" ("he" or "his") can be a "woman," whether or not a "male person" can be a "lawyer," whether or not a "client" can be a "lawyer," and whether or not a "client" can be a "woman," in particular one who hires a lawyer. In anaphoric reference, concepts can be identified for precisely the same set of cases that are identifiable in Meyer's (1970) task, when the sentence starts with "Some." In other words, concepts are anaphorically identifiable whenever the concepts overlap in Venn diagram terms as do thrones and chairs, or mothers and writers. In a sense, then, Meyer is studying the question of how people make anaphoric references.

Deeper into comprehension, concept identifiability plays an even larger role, especially in dealing with novelty. It is one of the basic processes that allows people to construct an interpretation of a new idea out of pieces of old ideas they have stored in memory. To take an example that may have some novelty left in it, suppose a person hears sentence (2) for the first time:

2. Dumbo the elephant could fly.

To make sense of this a listener will identify Dumbo the elephant with the first flying thing he can think of. This prob-

ably will be a bird rather than, say, a plane or a blimp, because in starting from "elephant" and "flying," the shortest applicable path usually will go through "bird" (though not for lovers of flying squirrels). Since having wings is crucial to a bird's flying, the listener will probably provide some sort of bird-like wings for Dumbo. The wings will come out of the sides somewhere, probably from the shoulders because they correspond to where birds' wings are located. In fact, by Disney's design, Dumbo used his very large ears as wings. This example illustrates that comprehension of new concepts is often based on identifying them by way of analogy with old concepts. The comprehension process in this case rests upon the identifiability of the concepts "elephant" and "bird." The process by which "elephant" is identified with "bird," we think is the same process we described in the previous sections.

There are other examples of how comprehension often involves identifying two concepts by way of analogy. Metaphor is one example, as shown in (3).

3. The boy's brother is a hippopotamus.

A person can decide by the rules of anaphoric reference that a person's "brother" and a "hippopotamus" are not the same thing. Because the metaphoric reference equates them, it is necessary to identify them analogically. Thus, just as Dumbo the elephant was given wings which are the most applicable property for flying that birds have, so for this case will the brother be given the most applicable properties of hippopotami. One can infer that his brother must be a large and languorous sort of chap. In fact, with the earlier example of Dumbo, the elephant's flying was treated essentially as a metaphor. Metaphor is really just the case of identifying two concepts that are not identifiable anaphorically.

But the process of identifying concepts with related concepts in comprehension is not restricted to farfetched examples like flying elephants and hippopotamus brothers. It occurs in everyday language in various ways, as seen in (4) and (5):

4. An old knob was fastened on the gate.

5. He hung his coat on the freshly painted door.

In dealing with (4), a person may never have seen or heard of a knob on a gate, but he can readily identify the gate with a door, or the knob with a latch on a gate to make sense of the sentence. In (5), there is an even more common use of identification of concepts. The reader can identify the door in the sentence with a door in memory (either a specific door or a composite door) that he has seen a coat on. The door in memory that he identifies it with need not have been freshly painted, nor need he ever have seen a freshly painted door. He can apply the paint in the same way he has seen it applied on other objects. Thus by identifying the door in the sentence with both a door in memory and a freshly painted object in memory, a person can construct an interpretation of a situation he has never witnessed.

In this last example, we have slipped back to the case where comprehension is a matter of identifying new concepts with old concepts that have the same name. But treating an elephant like a bird, or a brother like a hippopotamus are not very different from treating a gate like a door or one door like another door. The first two cases are only farther fetched. In other words, the semantic distance is greater between an elephant and a bird than between one door and another.

The problem of concept identifiability comes up in question answering as well as comprehension. Suppose a person is asked a question like (6) or (7):

- (6) Does a leopard have stripes?
- (7) Can a canary quack?

For questions like (6), where there is a similar concept with the property mentioned (in this case tiger), we have found (Collins and Quillian, in press) that people take a relatively long time to decide the answer is "no." This result is understandable because subjects tell us they think of the fact that a leopard is similar to a tiger which does have stripes. Our structural assumptions make the basis for the difficulty fairly clear. While people learn and store that tigers have stripes and leopards have spots, we doubt that they would ever learn or store that tigers do not have spots or that leopards do not have stripes. Thus, we would expect the fact that a tiger has stripes to be much more accessible starting from leopard than the fact that a leopard does not, inasmuch as the latter involves an inference: leopards have spots, and spots are not stripes. We will examine the nature of this inference further with example (7).

For (7), a person is likely to tell you the answer is no because it is ducks that quack and canaries aren't ducks. If so, the person has explicitly considered the question of whether canaries and ducks can be the same thing. We would argue that probably he has also tacitly considered whether singing, which is the sound canaries make and quacking are the same. Thus, the question of concept identifiability comes up twice in both (6) and (7).

In this case, the reasoning involved in the above inference is quite illogical, even though it sounds plausible. Just because a seal is not a dog is no reason to conclude that a seal doesn't bark. Similarly, just because a dog can bark does not mean that it can't howl or whine. In other words, maybe canaries quack when they aren't singing even though they are not ducks. The inference that spots contradict stripes was of the same type, but slightly more logical in that there are no animals that have both stripes and spots. The illogic of the inference used to decide that a canary doesn't quack could be made logical if people had the information stored that only ducks can quack. But that cannot be the case, since a person will be in doubt as to whether a goose can quack unless or until he realizes that a goose honks. The thing about geese is that they are so similar (i.e., have so many properties in common) to ducks that they might just quack if one isn't careful. But, if the person knew that only ducks quack, there should be no hesitation about geese once it is decided that geese aren't ducks. As another example, suppose people have stored that only dogs and seals bark. If they do then, of course, they can infer that wolves do not bark, since wolves are neither dogs nor seals. We think people who know how similar wolves and dogs are will be in doubt, unless, of course, they have heard of a wolf barking. A wolf is similar enough to a dog that, even if one knows they howl, there is no reason to suppose that they do not bark.

These examples suggest that there is a tradeoff in people's use of this kind of inference. For the case of quacking, the more different an animal is from a duck (e.g., a kangaroo), the more willing people are to conclude that the animal can't quack without knowing anything about the sound the animal makes.

The reason why knowledge about the properties of the animal in question enters into this tradeoff probably is based on a principle such as: "The more I know about the animal, the more likely I would know about its quacking if it did in fact quack." In any case, this kind of inference is basic to the way people answer questions, and we fear that computers will have to give up their past insistence upon rigor, if they want to be able to answer the range of questions people can answer.

There is another way that people obtain inferential power from this process of identifying a concept with a similar concept in question answering. Suppose a questioner needs information about the cost to ship tea to Boston from England, or information about the estimated number of schools in Boston. When the direct information on such questions is lacking for Boston but available for other cities, people often rely upon a type of analogical inference for an answer. They identify a city that is like Boston in the relevant respects and then infer the answer for Boston from the answer for the other city. For the question of costs of shipping tea to Boston, the strategy is to pick a nearby city (say New York or Providence) for which a figure is available. For the question about the number of schools in Boston, the strategy is to pick a city the same size as Boston, (say St. Louis) for which a figure is available. Then the answer for the other city may be adjusted to accommodate any difference between Boston and the other city on the relevant dimension. As Copi (1961) points out, analogy is at the basis of most of our ordinary inferences from past experience into the future, as when one reads a book because he enjoyed the author's previous books. But we would argue further that analogy, in fact, underlies every aspect of our inferential reasoning.

In this section, we have enumerated a number of different examples from natural language where the processing involved identifying one concept with a similar concept or, on the other hand, finding a basis for distinguishing one concept from a similar concept. To psychologists, this is an old process in a new guise. From the first point of view, it is the process of generalization; from the other point of view, the process of discrimination. Surely, Pavlov would think it a great joke to find us caught in this web.

Knowing the new uses that evolution makes of old organs, it is not too surprising to find such a primitive capability put to heavy use in the most sophisticated of man's talents. This may be an area where linguists might utilize psychological knowledge for a change instead of the other way round. As Brown (1970) points out, this would probably be the first time in the history of psycholinguistics.

Generalization and discrimination do not come naturally to present-day serial computers, but by using tags to simulate spreading activation, we think it is possible to develop generalization and discrimination processes in computers. The strategies outlined in the last section were examples of such processes.

3.6 Two Uses of Imagery

This work started out by investigating structure and processing in a semantic network (Quillian, 1968, 1969) of interconnected concepts. Imagery was happily ignored in order to keep the problems down to a manageable size, though there was an unspoken assumption that images could be generated from concepts.

In this section, we will not try to deal with imagery systematically, but only mention where it fits in the scheme of computer-language processing. Imagery has intruded into our experimental studies of question answering because subjects usually give some reason based on imagery for rejecting false sentences (Collins and Quillian, in press). While imagery probably has many uses, we want to mention two that are particularly relevant to the relation between imagery and concepts.

When imagery intruded into our subjects' reports in rejecting false sentences, it turned out to be used in a way that Quillian had developed to handle the problems of anaphoric reference within his semantic network. In other words, the strategy that people reported using with images was the same strategy Quillian planned to use with semantic properties. In particular, the strategy was to reject the possibility of identifying two concepts whenever the concepts had properties where the attributes (or relations) matched and the values were different. In the anaphoric-reference example of the last section (1), "he" did not refer to "woman," because, even though they have a common superset ("person"), one has the value "male" and the other "female" on the attribute "sex." The use of this strategy with sensory properties is exactly the same. To take one example, a subject rejected the identifiability of a pearl and a bean after comparing them in imagery. He noticed they appeared to be the same size, but that they are different in shape and color, which formed the basis for his rejection. Imagery may be an efficient way to compare concepts to find a mismatch, because images can be manipulated in a way that properties cannot. Perhaps, interactions between separately stored concepts and their properties can be evaluated more readily if the concepts are generated together in imagery. Such an interaction can be illustrated by an example from an earlier paper (Collins and Quillian, in press). Several subjects reported that they rejected "A limousine has a rudder" as false by imaging the rudder on the back of a limousine. It is an imagined interaction that produces a mismatch with memory.

The second aspect of imagery we want to mention makes it seem as if there is much more information in internally generated images than there need be. We would argue that images (not eidetic images, however) contain much less information than pictures. What appears to be richness in images, we think, derives from changing images by changing the concepts from which they are generated. For example, if a person thinks back in memory as to where he left his keys in his office, the procedure might be something like the following:

- (1) Generate an image of the office and pick the most likely piece of furniture there (say a desk) where the keys might be.
- (2) Generate an image of the desk top, and scan that for the keys.
- (3) If not found, then generate the image of the inside of the desk drawer and scan that for the keys.
- (4) If still not found, regenerate the image of the office and pick the next most likely piece of furniture.
- (5) etc.

The ability to change from imaging one concept to imaging another would give a computer system the kind of power that a zoom lens on a camera provides. The difference is that, unlike the camera

picture, not all the information about the desk top and the drawer, obtained by changing images, is in the original image of the office. This kind of manipulation of internally-generated images from a concept network is still only a gleam in some computer-mind's eye.⁶

3.7 Induction and Learning

How information is put into memory obviously has much to do with how it can be retrieved. Until now we have only dealt with inferences in retrieving information as opposed to inferences in storing information. Generally, the inferences made in the process of storing information are referred to as inductions. For example, if a child sees red flowers on several occasions that people refer to as roses, he may well induce from this experience the information that roses are red (i.e., that particular pinkish-red most roses are) as a property which can be evoked by the concept "rose." The fact that the property of redness comes to be evoked by the concept of rose is non-committal as to whether all roses are red, most roses are red, some roses are red, or only a few roses are red (though in the latter case redness probably would never come to be evoked). To evaluate whether all, or most, or some, or a few roses are red, given that a person has not learned the answer directly, would require a search of memory starting at "color" and "rose" to find connections through colors other than red. Depending on how many non-red instances of roses are found (with a given amount of effort or as compared to red ones), some estimate about the correct answer can be made.

The importance of induction to language processing lies in the fact that most properties of concepts are not learned directly, but are derived from specific instances of the concept. A person probably stores the fact that birds have wings, because he sees many instances of birds with wings. Once he stores by induction from cases that birds have wings, he can then deduce that a new instance he knows to be a bird, say a mudlark, must therefore have wings.

The fundamental question about induction is how to decide which properties of specific instances should be stored with the superset? The problem has been attacked for computers by Becker (1970) and Winston (1970). Both deal with the problem in terms of generalization and discrimination; apparently, Pavlov has ensnared them as well. Both their systems are designed to learn concepts from graph-structure descriptions of visual scenes, and so both deal with the problem in terms of visual properties rather than symbolic (i.e., non-imageable) properties. If concept learning is in terms of verbal inputs, as might be the case with "vehicle," then the problem would arise in terms of symbolic properties. For instance, if a person is told that cars, buses, and trains are all vehicles, then he may induce from these instances which of their properties apply to the concept vehicle. Again, we would argue that the problem is fundamentally the same whether posed in visual terms or symbolic terms, though visual concept learning is prior developmentally. We will outline the process (for more detail, see Becker or Winston) with an example from the world of birds where we feel relatively safe from flights of fancy.

Suppose one sees a cute little red feathery creature with wings and a beak, etc. and that it is referred to as a "bird." Next maybe a larger (by a factor of three) black feathery creature with wings and a beak that is referred to as a "bird." Next maybe a little yellowish feathery creature with wings and a beak that is referred to as a "bird." Each time a person hears a repetition of the word "bird" in the sequence, it leads him to those properties he had stored previously with bird. If a semantic search starts with "bird" on the one hand and the instance he is looking at on the other, then the connections found will be through properties that the bird in memory and the bird at hand have in common. Therefore, it is precisely those common properties that will be made more accessible (i.e., wings, feathers and beaks) and the others will not (i.e., color and size).⁹ The process as described is pure generalization. Winston (1970) points out that when learning "arch" in this way, the top piece of the arch may be a rectangle one time and a wedge the next time. The strategy he uses is to assume that, when such a difference is found, the top piece can be any instance of the lowest common superset (in this case, block) of the top pieces seen so far. This is exactly the way we would have it.

Suppose instead of each instance being referred to as a "bird" in the preceding sequence, the first was referred to as a "cardinal," the second a "vulture," and the third a "canary." When a person sees the canary in this sequence, he will search his memory for a concept that has the properties he sees the canary to have. If we assume he does not have the concept "canary," then he will intersect with the concept in memory most like a canary; i.e., the one with the most properties in

common with the canary he sees. If he has already formed the concept for bird, then he will find the concept "bird." If he only had learned "cardinal" and "vulture" then he would probably find the concept "cardinal," because it is more like a canary than a vulture is. Thinking of the concept "bird" or "cardinal" when presented with a canary is like thinking of a tiger when asked "Does a leopard have stripes"? [Example (6), Section 3.5]. Whether he locates the concept "bird" or "cardinal" in his search, the name will be different from the name "canary" given to the thing he is looking at. Because the names are different, this forces him to locate those properties that are different (e.g., color). It is the distinguishing properties that are then stored with the new concept "canary." In Winston's terms a bird or cardinal is a "near miss" for the concept "canary." As we described it in this example, the process is pure discrimination learning. Generally though, a person would see several positive instances of canaries in forming the concept "canary." Hence, the course of true learning will be both a generalization and discrimination process. But in learning the concept "bird," we doubt that there are any of the "near misses" that Winston's program seems to rely on. That is to say, people don't look at bats in order to learn the concept "bird."

There are several non-obvious implications of this view. One important implication is that higher-level concepts¹⁰ like bird are learned mostly by generalization and lower-level concepts like canary mostly by discrimination. Of course, to the degree that people relate birds to other animals in learning the concept "bird," then obviously they will have to discriminate birds from the other animals. But, if one considers the relative similarity between birds and dogs, say, as opposed to the

similarity between canaries and cardinals, it should be clear why we think that learning the concept "bird" will involve much less discrimination than will learning the concept "canary." And, if one considers the relative variability in properties among birds as opposed to the variability among canaries, it should be clear why we think that learning the concept "bird" will involve more generalization than will learning the concept "canary." Anglin (1970) in a recent monograph conducted several experiments with children to find out whether the learning of word meanings was a generalization process or a discrimination process. He concluded that it was a generalization process but the lowest-level concepts he used were boy and girl. In our view, his selection of words would prevent his finding the discrimination learning that takes place at lower levels, such as between canaries and cardinals.

It turns out that, when negative properties are stored in memory (such as "Penguins can't fly" or "A vest doesn't have sleeves, in Table 1), it is because learning depends on discrimination. Hence, there is always some similar concept (a confusable concept) which has the given property. In retrieval, the confusability will slow a person down so that it will be difficult to decide, for example, "Does a vest have sleeves"? That is to say, in starting at "vest" and "sleeves," the semantic search is likely to turn up "suit jacket" or some other similar concept with sleeves. Because a concept like "bird" is learned mostly by generalization, there are not likely to be any negative properties stored with it, e.g., "Birds do not have antlers."

As Becker points out, properties never become defining properties by either generalization or discrimination. This is the general semanticist's old argument (Korzybski, 1933) that there is no essence of "chairness" or "pencilness." To take Winston's example of arches, people can still identify Roman or Greek arches as arches, even though the top piece has been knocked down and is lying on the ground. The lack of defining properties in this world makes life much harder for computers.

As the process of learning the concept "canary" was described, it was mostly a discrimination process but partly a generalization process. To the degree it was a discrimination process, the fact that canaries have wings and feathers would be ignored, since wings and feathers do not differentiate canaries from cardinals or birds. However, to the degree that it is a generalization process, wings and feathers would always co-occur with instances of canaries, and hence would be stored as properties of canaries. Then, the question arises as to why our data show (as argued in section 2.5) that people decide about a sentence like "A canary has wings" by inference from the fact that a canary is a bird and birds have wings. We think the finding has two bases.

First, for many kinds of birds (and other things as well) most people do not have stored much more than the fact that they are birds. This is because a person does not form concepts like mudlark by a process of generalization and discrimination. Instead, he may be told or he may infer from something he reads, that a mudlark is a bird. When some psychologist asks him to decide if "A mudlark has wings" is true or false, he responds

true, because he knows a mudlark is a bird. In constructing sentences for our study (Collins and Quillian, 1969), we made an effort to choose instances (e.g., wren) in such a way that the superset property (e.g., has wings) was not particularly associated with the instance. Hence, by design the sentences used were ones likely to be decided by inference. There is nothing in the theory, however, that prevents storing superset properties with particular instances, and we certainly think it is a common practice.

Second, even if a superset property is stored directly with an instance, it may be faster to retrieve it by inference from the superset. This might happen if the property is fairly inaccessible from the instance, but highly accessible from the superset. The learning process we described would tend to produce such a difference in accessibility, because generalization of the property (having wings) for the superset (bird) makes the property more accessible, and discrimination for the instance (canary) does not make the property more accessible. As we argued earlier, learning the superset is mostly a generalization process and learning the instance is mostly a discrimination process. Similar differences in accessibility can produce a wrong response to the question, "Can a penguin fly"? If the person thinks first that a penguin is a bird, and birds fly, before he retrieves the fact that penguins can't fly, he may well give an incorrect "yes" response to the question.

One final comment about generalization and discrimination in computers. In Section 3.5, we talked about generalization and discrimination in comprehension and retrieval. There it was necessary to identify a concept with a similar concept in order

to comprehend a sentence or answer a question. In this section, we have been discussing generalization and discrimination in learning. Here it is necessary to identify the object at hand with a similar concept in memory. Whether the concepts that are identified in this way are treated similarly (generalization) or differently (discrimination) is imposed in either case not by the nature of the concepts identified but by the constraints of the task. In our example of metaphor, the hippopotamus and the boy's brother are treated similarly by the way the sentence forces them to be identified. In generalization learning, a cardinal and a canary are treated similarly because they are both called birds. In discrimination learning, on the other hand, a canary would be distinguished from a cardinal, and in anaphoric reference a boy's brother would not be identified with a hippopotamus. The importance of generalization and discrimination to learning, comprehension and retrieval makes it imperative, we think, that these processes be treated by any computer implementation in the same terms.

3.8 Analogy as a Linguistic Construct

Linguistic approaches to semantic theory, which are at least partially reflected in Kintsch's (1972) paper in this volume, typically specify human semantic knowledge in terms of selection restrictions. Therein lies a major difference between our proposed theory and linguistic-based theories. The difference can be illustrated with an example from Kintsch's paper where he points out that it is permissible to say "The child grew," "The corn grew," and "The farmer grew the corn," but not "The parents grew the child." He explains this in terms of a selection restriction on the use of grow that prohibits a human or animal object in the presence of an agent. Agent and object in this description are cases, which are special kinds of relations used in describing verb concepts. In our thinking, cases function for

verb concepts much like inference-carrying relations function for noun concepts, in that they are used extensively in language processing, and there are a small number of them that are used quite frequently. On this we are in agreement with Kintsch.

The disagreement can be illustrated by a setting from one of the Oz books by Baum (1908). In the story, Dorothy falls into the earth during an earthquake and lands in a city where apparently there are no children among the people. As the story progresses, she finds out that the people are vegetables and that they have a special garden just outside the city. It turns out that the adults grow children in the garden until they are ripe and then the adults pick them from the garden to become members of the community. Now, in this setting, talk of adults growing children sounds quite natural; and, of course, it is because the people are vegetables and it is quite easy to analogically identify vegetable children with plants. Even without such an elaborate plot, it is perfectly easy to understand what is meant by "A farmer grew sheep on his ranch" or "The parents grew their child in isolation."

Everybody has heard of people growing corn, or flowers, or grass; and from these examples, a person might induce the concept of plant as the object of grow in the presence of an agent. Then, when a person hears that a lazy man grows rocks in his garden or a scientist grows a theory during his coffee breaks, these can be understood by analogy with how people grow plants. Kintsch postulates special metaphor rules to deal with such cases, but in our view that merely sets up two processes where only one pro-

cess is necessary. In comprehending such sentences, we think people will always try to find a meaning by identifying the object of grow in the presence of an agent with the concept plant, or whatever object concepts are learned or induced for the verb grow. The more far-fetched is the analogy, or in other words the more distant is the object concept from plant in semantic memory, the odder the sentence will sound; i.e., the more semantically anomalous it will be. But the point is that there need be no switch or line on one side of which is semantic acceptability and on the other side of which is metaphor or anomaly.

In our view, there is a continuum from semantic acceptability to metaphor to anomaly depending on how removed the given object (e.g., trees, horses, children, rocks) is from plant. If you ask a person whether a particular sentence is a metaphor or not, he can answer; but we think he must specifically evaluate the sentence to do so. The test he would probably apply is whether the given object is acceptable as an instance of plant (i.e., whether there is any overlap in Venn diagram terms between the given object and plants). If it is, then he will say it is semantically acceptable and if not he will call it a metaphor. But if you ask a person whether "He grew oysters in his pond" is a metaphor or not, he is likely to have trouble seeing that it is, because oysters, though animals, are so plant-like. The fact that there are hard cases to decide argues quite strongly that there is a continuum rather than a switch.

It is possible to think about selection restrictions in terms of storing negative information rather than positive information in memory. For example, one might store the information that grow prohibits a human or animal object in the presence of an agent, rather than the information that grow takes plants as the object in the presence of an agent. We are not sure whether Kintsch thinks that negative information is stored or not. As we argue in the preceding section, it seems to us that negative information is only stored in the process of discrimination learning. There is nothing about the process of hearing sentences where people grow flowers, or corn, or grass that would produce discrimination learning. Hearing such sentences would be a case of pure generalization learning, unless children learn that saying something like "The farmer grew the horses" or "The mother grew children" is wrong. We seriously doubt that they are told such sentences are wrong. Should they utter such sentences (if they do at all), we doubt that they would ever be corrected for doing so. But unless they learn such sentences are wrong, our position means that children would never form the negative property that animal or person is prohibited as the object of grow in the presence of an agent. They would only learn the positive property that plants are often the object of grow in the presence of an agent.

One aspect of the learning process is worth pointing out. Suppose all a child knows about growing is that children grow up. Then, if he hears the sentence "The corn grows in the field" he can understand it by identifying corn with children. But this use of corn is a metaphor for the child. Suppose later he hears that "flowers grow." He is more likely to relate the

flowers to corn than children, and so the child is likely to understand the flowers growing by analogy with corn growing. This too is a metaphor for the child, because clearly flowers are not corn. In general then, depending on what one has heard before, one man's semantically acceptable sentence is another man's metaphor.

4. THE THEORY AND EXPERIMENTAL PSYCHOLOGY

The theory we have outlined is not designed specifically to yield clear-cut predictions about experiments. Nor did it arise as an explanation of results found in experiments. It started out as a strategy for dealing with language in a computer, and much of the shape it is growing derives not from watching people perform, but instead from what programmers figure out about people when they try to build parts of a language-using machine. The theory's main function is to provide guidelines for how to go about building a memory. Hence, it contains descriptions of internal structures and processes, rather than input-output transfer functions for various stages of processing. The processes we are hypothesizing can be put together in too many ways, depending on the strategies of the subject, to always yield output predictions from various input conditions. However, we have tried to show that it leads to some hypotheses that can be tested experimentally.

We think psychology can profit from trying to build a language-using machine, just as the theory of flying has profited from trying to build a flying machine. Think of all the useless experiments on flying that could have been done with birds.

5. REFERENCES

Anglin, J. M. The growth of word meaning. Research Monograph No. 53, Cambridge, Mass.: M.I.T. Press, 1970.

Baum, L. F. Dorothy and the wizard in Oz. Chicago: Reilly and Lee Co., 1908.

Baylor, G.W. A treatise on the mind's eye. Institut de Psychologie, Université de Montreal, 1971.

Becker, J. D. The modelling of simple analogic and inductive processes in a semantic memory system. In the Proceedings of the International Joint Conference on Artificial Intelligence, D. Walker and L. Norton (Eds.), 1969.

Becker, J. D. An information-processing model of intermediate-level cognition. Ph.D dissertation, Stanford Artificial Intelligence Project, 1970.

Begg, I. & Paivio, A. Concreteness and imagery in sentence meaning. Journal of Verbal Learning and Verbal Behavior, 1969, 8, 821-827.

Bobrow, D. G. and Fraser, J. B. An augmented state transition network analysis procedure. In the Proceedings of the International Joint Conference on Artificial Intelligence, D. Walker and L. Norton (Eds.), 1969.

Bobrow, D. G., Fraser, J. B., & Quillian, M. R. Automated language processing. In the Annual Review of Information Science and Technology, Vol. 2, C. A. Cuadra (Ed.), New York: Wiley, 1967.

Bower, G. R. Mental imagery and associative learning. In Cognition in Learning and Memory, L. W. Gregg (Ed.), New York: Wiley, in press.

Brown, R. Psycholinguistics. New York: Free Press, 1970.

Bruner, J. S. & Minturn, A. L. Perceptual identification and perceptual organization. Journal of Genetic Psychology, 1955, 53, 21-28.

Carboneil, J. R. Mixed-initiative man-computer instructional dialogues. Ph.D dissertation, M.I.T., 1970. Also BBN Report No. 1971. Reprinted in part as: AI in CAI: An artificial intelligence approach to computer-aided instruction. IEEE Transactions on Man-Machine Systems, 1970, MMS-11, 190-202.

Chase, W. G. & Clark, H. H. Mental operations in the comparison of sentences and pictures. In Cognition in Learning and Memory, L. W. Gregg (Ed.), New York: Wiley, in press.

Collins, A. M. and Quillian, M. R. Retrieval time from semantic memory. Journal of Verbal Learning and Verbal Behavior, 1969, 8, 240-247.

Collins, A. M. and Quillian, M. R. Facilitating retrieval from semantic memory: The effect of repeating part of an inference. Acta Psychologica, Attention and Performance III (A. F. Sanders, ed.) 1970a, 33, 304-314.

Collins, A. M. and Quillian, M. R. Does category size affect categorization time? Journal of Verbal Learning and Verbal Behavior, 1970b, 9, 432-438.

Collins, A. M. & Quillian, M. R. Experiments on semantic memory and language comprehension. In Cognition in Learning and Memory, L. W. Gregg (Ed.), New York: Wiley, in press.

Collins, A. M. and Quillian, M. R. Category and subcategories in semantic memory. In preparation.

Copi, I. M. Introduction to logic. New York: Macmillan Company, 1961.

Gelernter, H. Realization of a geometry-theorem proving machine. In Computers and Thought, E. A. Feigenbaum & J. Feldman (Eds.), New York: McGraw-Hill, 1963.

Giuliano, V. E. Analog networks for word association. IEEE Transactions on Military Electronics, 1963, MIL-7, 221-225. Reprinted in part in: Memory and Attention by D. A. Norman, New York: Wiley, 1969.

Johnson-Laird, P. N. & Stevenson, R. Memory for syntax. Nature, 1970, 227, 412-413.

Kintsch, W. Notes on the structure of semantic memory. To appear in Organization and Memory (ed. by E. Tulving and W. Donaldson) New York: Academic Press, 1972.

Kolers, P. A. Interlingual facilitation of short-term memory. Journal of Verbal Learning and Verbal Behavior, 1966, 5, 314-319.

Korzybski, A. Science and Sanity. Lancaster, Pa.: Science Press, 1933.

Landauer, T. K. & Freedman, J. L. Information retrieval from long-term memory: Category size and recognition time. Journal of Verbal Learning and Verbal Behavior, 1968, 7, 291-295.

Meyer, D. E. On the representation and retrieval of stored semantic information. Cognitive Psychology, 1970, 1, 242-300.

Olney, J. C. Some patterns observed in the contextual specialization of word senses. Information Storage and Retrieval, 1964, 2, 79-101.

Paivio, A. Mental imagery in associative learning and memory. Psychological Review, 1969, 76, 241-263.

Polyani, M. The tacit dimension. Garden City, New York: Doubleday, 1966.

Postman, L. A pragmatic view of organization theory. To appear in Organization and Memory (ed. by E. Tulving and W. Donaldson) New York: Academic Press, 1972.

Quillian, M. R. The teachable language comprehender: A simulation program and theory of language. Communications of the ACM, 1969, 12, 459-476.

Quillian, M. R. Semantic memory. In Semantic Information Processing, M. Minsky (Ed.), Cambridge, Mass.: M.I.T. Press, 1968.

Rubenstein, H., Lewis, S. S., & Rubenstein, M. A. Homographic entries in the internal lexicon: Effects of systematicity and relative frequency of meanings. Journal of Verbal Learning and Verbal Behavior, 1971, 10, 57-62.

Schaeffer, B. & Wallace, R. Semantic similarity and the comparison of word meanings. Journal of Experimental Psychology, 1969, 82, 343-346.

Schaeffer, B. & Wallace, R. The comparison of word meanings. Journal of Experimental Psychology, 1970, 86, 144-152.

Shiffrin, R. M. Forgetting: Trace erosion or retrieval failure? Science, 1970, 168, 1601-1603.

Simmons, R. F., Burger, J. F., & Long, R. E. An approach toward answering English questions from text. AFIPS Conference Proceedings, Vol. 29, Washington, D.C.: Spartan Books, 1966.

Simon, H. A. What is visual imagery? An information processing interpretation. In Cognition in Learning and Memory, L. W. Gregg (Ed.), New York: Wiley, in press.

Thorne, J. P., Bratley, P., & Dewar, H. The syntactic analysis of English by machine. In Machine Intelligence 3, D. Michie (Ed.), New York: American Elsevier, 1968.

Wilkins, A. J. Categorization time and category size. Journal Verbal Learning and Verbal Behavior, in press.

Winston, P. H. Learning structural descriptions from examples. Ph.D thesis, M.I.T., Project MAC, MAC TR-76, 1970.

Wittgenstein, L. Philosophical Investigations. (Translated by G. E. M. Anscombe) Oxford: Blackwell, 1953.

Woods, W. A. Context-sensitive parsing. Communications of the ACM, 1970, 13, 437-445.

FOOTNOTES

1. This research was supported by the Advanced Research Projects Agency of the Department of Defense and was monitored by the Air Force Office of Scientific Research under Contract No. F44620-67-C-0033.

We should like to thank the Artificial Intelligentsia at Bolt Beranek and Newman Inc. whose ideas we may have surreptitiously stolen; namely, Joseph D. Becker, Anthony Bell, Daniel G. Bobrow, Jaime R. Carbonell, and William A. Woods. We also want to thank David Meyer and Benson Schaeffer for their comments on an earlier draft of this paper.

2. This is not say that computers may not in the future function in different terms.
3. Properties are made up of two parts: relations and objects. For adjectival properties, relations are called attributes and objects are called values. Red and square are typical values, and they refer to concepts. Usually red is said to be a value of the attribute color. We are arguing that, in general, both attributes and values are more complex than this. Thus, for the concept "canary" the attribute color may have embedded the fact that it is only on the surface, and the value yellow may have embedded that it is light yellow, or one of the several yellows that people can distinguish in memory.

4. The conceptual identity of buying and selling could be tested by a modification of Koler's (1966) technique. He showed that presenting a word and its translation at different places in recall lists given to bilingual subjects improved recall of the word as much as presenting the same word twice. If buy and sell refer to the same concept, they should reinforce each other in a similar way. However, a recent study of Johnson-Laird and Stevenson (1970) suggests that the two words must refer to the same concept.
5. This, we think is partly the reason why our results (Collins and Quillian, 1969) in a true-false reaction time task show people to be faster in deciding that "A canary is a bird," for example than they are in deciding that "A canary can sing." We assume that both the superset "bird" and the property "can sing" are stored directly with canary in most people's memory.
6. This raises the problem of whether all properties should be stored as supersets, since this makes inferences easier. It is easier to retrieve the properties of military dictatorships (e.g., they imprison dissidents) for Paraguay, if Paraguay is stored as a military dictatorship rather than as having a military dictator. Similarly, it would be easier to retrieve properties of hot objects (e.g., they burn hands) for an oven rack, if oven rack is stored as a hot object rather than as being hot.

We don't know where to draw the line in setting up supersets, but it seems a bit much to have supersets such as hot objects or objects on Gorky Street. More likely a person infers that oven racks will burn hands by analogy with the fact that irons (or whatever object he has learned about) burn hands. In the same sense though, canine and marsupial may not generally be supersets either, except to zoologists. If asked whether marsupials hop, people who know kangaroos are marsupials will probably answer "yes," but they must get the property from kangaroos and not from marsupials. That is to say the property is inferred from an example just as with the iron.

7. There is a complication here that will be discussed in relation to induction (see Section 3.7).
8. Such a capability for maps is now being developed by Jaime Carbonell at BBN.
9. Increasing accessibility will also produce forgetting. As some properties become more and more accessible, the likelihood of retrieving properties that do not becomes smaller and smaller. This is in line with Shiffrin's (1970) recent finding which suggests that forgetting is purely a failure in retrieval.
10. Of course, there are probably exceptions such as mammal. Mammals, though a higher-level concept may be learned mostly by discrimination from reptiles, amphibian, fish, and birds.